

04-07-06

A

UTILITY PATENT APPLICATION TRANSMITTAL

Submit an original and a duplicate for fee processing
(Only for new nonprovisional applications under 37 CFR §1.53(b))

ADDRESS TO:

Assistant Commissioner for Patents
Box Patent Application
Washington, D.C. 20231

Attorney Docket No. 202732

First Named Inventor Colleran

Express Mail No. EL304646420US

JC678 U.S. PTO

09/543930

04/06/00

APPLICATION ELEMENTS

1. ☒ Utility Transmittal Form
2. ☒ Specification (including claims and abstract) [Total Pages 34]
3. ☒ Drawings [Total Sheets 10]
4. ☒ Combined Declaration and Power of Attorney [Total Pages 4]
 - a. ☒ Newly executed
 - b. ☐ Copy from prior application

[Note Box 5 below]

 - i. ☐ Deletion of Inventor(s) Signed statement attached deleting inventor(s) named in the prior application
5. ☐ Incorporation by Reference: The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. ☐ Microfiche Computer Program
7. ☐ Nucleotide and/or Amino Acid Sequence Submission
 - a. ☐ Computer Readable Copy
 - b. ☐ Paper Copy
 - c. ☐ Statement verifying above copies

ACCOMPANYING APPLICATION PARTS

8. ☒ Assignment Papers (cover sheet and document(s))
9. ☐ Power of Attorney
10. ☐ English Translation Document (if applicable)
11. ☐ Information Disclosure Statement (IDS)
 - ☐ Form PTO-1449
 - ☐ Copies of References
12. ☐ Preliminary Amendment
13. ☒ Return Receipt Postcard (Should be specifically itemized)
14. ☐ Small Entity Statement(s)
 - ☐ Enclosed
 - ☐ Statement filed in prior application; status still proper and desired
15. ☐ Certified Copy of Priority Document(s)
16. ☒ Other: Check in the amount of \$1156.00

17. If a **CONTINUING APPLICATION**, check appropriate box and supply the requisite information in (a) and (b) below:

(a) ☐ Continuation ☐ Divisional ☐ Continuation-in-part of prior application Serial No. _____
Prior application information: Examiner _____; Group Art Unit: _____

(b) Preliminary Amendment: Relate Back - 35 USC §120. The Commissioner is requested to amend the specification by inserting the following sentence before the first line:

"This is a ☐ continuation ☐ divisional of copending application(s)

☐ Serial No. _____, filed on _____

☐ International Application, filed on _____, and which designates the U.S."

APPLICATION FEES

APPLICATION FEES				
BASIC FEE				\$690.00
CLAIMS	NUMBER FILED	NUMBER EXTRA	RATE	
Total Claims	35 -20=	15	x \$18.00	\$270.00
Independent Claims	5 -3=	2	x \$78.00	\$156.00
<input type="checkbox"/> Multiple Dependent Claims(s) if applicable			+\$260.00	\$
Total of above calculations =				\$1116.00
Reduction by 50% for filing by small entity =				\$()
<input checked="" type="checkbox"/> Assignment fee if applicable			+\$40.00	\$40.00
TOTAL =				\$1156.00

UTILITY PATENT APPLICATION TRANSMITTAL

Attorney Docket No. 202732

19. ☐ Please charge my Deposit Account No. 12-1216 in the amount of \$.
20. ☒ A check in the amount of \$1156.00 is enclosed.
21. The Commissioner is hereby authorized to credit overpayments or charge any additional fees of the following types to Deposit Account No. 12-1216:
- a. ☒ Fees required under 37 CFR §1.16.
- b. ☒ Fees required under 37 CFR §1.17.
22. ☒ The Commissioner is hereby generally authorized under 37 CFR §1.136(a)(3) to treat any future reply in this or any related application filed pursuant to 37 CFR §1.53 requiring an extension of time as incorporating a request therefor, and the Commissioner is hereby specifically authorized to charge Deposit Account No. 12-1216 for any fee that may be due in connection with such a request for an extension of time.


23. CORRESPONDENCE ADDRESS

Rattan Nath, Registration No. 43,827
Leydig, Voit & Mayer, Ltd.
Two Prudential Plaza, Suite 4900
180 North Stetson
Chicago, Illinois 60601-6780
Telephone: (312) 616-5600
Facsimile: (312) 616-5700

Name

Rattan Nath, Reg No. 43,827

Signature




Date

April 6, 2000

Certificate of Mailing Under 37 CFR §1.10

I hereby certify that this Utility Patent Application Transmittal and all accompanying documents are being deposited with the United States Postal Service "Express Mail Post Office To Addressee" Service under 37 CFR §1.10 on the date indicated below and is addressed to: Assistant Commissioner for Patents, Box Patent Application, Washington, D.C. 20231.

Matthew W. Olson



April 6, 2000

Name of Person Signing

Signature

Date

RESPONSIVE USER INTERFACE TO MANAGE A NON-RESPONSIVE APPLICATION

5

TECHNICAL FIELD

This invention relates generally to development of computer operating systems for executing software applications and more particularly to the development of user interface capabilities provided to application programs by an operating system.

10

BACKGROUND OF THE INVENTION

There is an increasing appreciation for the need for aesthetics and ease of use in the design of user interfaces ("UI") to encourage use of computer applications and reduce frustration. To term the new user interfaces design strategies in cutting edge software as being "ergonomic" is not an overstatement. Such interfaces are a combination of functionality coded in the application and supported by the underlying platform, which may be an operating system ("OS"). Thus, the OS provides significant support for realizing an effective UI, and such support is an important factor in the competition between different OS available in the market.

Application programs are typically written for execution on a particular platform. Such a platform may be defined by a virtual machine offering a relatively invariant context regardless of the underlying hardware or an OS that is tied down to a particular hardware platform or even an application providing an execution context for another application, e.g., a plug-in module. The platform includes a plurality of services and appropriate management strategies to allow different applications access to system resources. These services provide many of the functions that applications are expected to

use and thus free the application writer from worrying about the more mundane implementation details such as security, reading and writing data, managing memory resources, and I/O functions. Not surprisingly, this is an effective strategy since the services provided by the OS are used repeatedly. Having a single coherent
5 implementation reduces complexity and enhances design of stable computing systems capable of executing different applications.

OS for computers, and for personal computers in particular, have undergone many changes. These changes not only improve the range of services available to applications, typically through an Application Programming Interface (API), but also provide new
10 architectures that reduce the likelihood of crashing the entire computer system due to errors including those ascribable to applications or users. An important advance in designing stable computing systems has been the development of multithreaded systems, which are described further below.

Traditional OSs for personal computers used a single threaded architecture in
15 which programming code was executed in a serial fashion. Any step or sequence of steps could result in a fatal bottleneck. A user is often hard-pressed to distinguish delays from a system failure. An improvement was a computing environment where if a busy application, as opposed to a non-functional application, could indicate that it was busy to the user, or preferably relinquish control to the OS to allow other applications and tasks to
20 execute before being granted control again. Such computing environments were at the mercy of a well-behaved and cooperative application.

In contrast, in a multithreaded architecture the OS exercises greater control over the execution of different tasks. The OS schedules slices of time on the processor for

identified units termed threads. A thread is a path of execution within a process that is recognized and provided time on the processor by the OS. Each application usually has at least one thread and, thus, is assigned time in accordance with its relative priority.

Additionally, often an application executes in its own address space, particularly if it is treated as a process by the OS. This results in isolation of different processes from each other and contributes to greater system stability and robustness even when executing faulty applications.

It should be understood that the term thread refers to code that is provided execution time slices by an external entity, e.g., the OS. This does not foreclose a developer of an application to define a path of execution within the application such that the application directly controls the time allocated to a particular path while the OS may be unaware of its existence. For clarity, such developer defined execution paths are referred to as “fibers” as opposed to threads. The benefit in using threads is that, even if an application is defective it cannot hold up rest of the computing system, since it never has exclusive control of the computing environment. However, this statement should not be interpreted to imply that rogue applications designed to subvert the OS are not possible in a multithreaded computing environment. Accordingly, for conceptual convenience, threads may be generalized to scheduled code paths.

At least one of the threads assigned to an application is termed its main thread and typically implements the UI for the application. Modern UIs are usually graphical interfaces formed by a plurality of elements termed window elements or just windows. These window elements, controlled by the threads in a system, are displayed to a user. The display space available on the desktop is referred to as “real estate.” The desktop is

typically implemented with the help of routines such as `paint()` to continually update it to reflect the current state and order of the windows displayed on it. A desktop thread, which is a system thread, manages painting of the desktop windows and also handles miscellaneous system events. Some windows that are fully or partially occluded by other windows may not be displayed in their entirety. The window in the foreground commands the attention of the user and has focus. Usually, the focus shifts as other windows are selected to be in the foreground.

Since window elements are not only output symbols but also aid in the input of events, each thread managing a window element is required to have a message loop. This loop allows messages to be sent for handling by the window element on the desktop. Similarly, messaging allows processing of input events by the OS, the application program or even the user by presenting them in an orderly fashion.

In addition, user input to the computer system is received from the hardware devices via a thread termed the raw input thread (RIT). The OS sorts the input events such as keystrokes, mouse movements, mouse clicks and the like in the RIT queue and forwards them to the appropriate application input thread queues or OS routines for handling. Input events are usually posted to an application owning the window element in focus. Furthermore, not all events are handled with the same priority. Some events, e.g., key combinations such as Alt-Tab keys shift focus from one window to another and are handled out of turn. Since each application has its own input queue, the failure of an application to handle its own queue entries is not fatal. In other words, it does not “hang” the entire system by making it unresponsive to further inputs.

The OS also provides implementation code for a variety of graphical objects and windows for use by application programs. This not only saves creators of applications the complex task of writing graphical code, but also promotes standardization of the UI.

An example of such an interface is the familiar window interface in the “WINDOWS®”

5 OSs manufactured by the “MICROSOFT®” Corporation of Redmond, Washington. The interface includes an enclosing rectangle, system buttons for closing, maximizing or restoring the window element, and an enclosed area that is available to the application/user to display arbitrary graphical or text symbols. The enclosing rectangle edges and corners are also usable for resizing the rectangle with the use of a pointing
10 device such as the familiar mouse. Other designs for UIs are also possible along with modifications to the familiar designs, e.g., by varying color schemes/designs, transparency and other properties.

As described previously, the desktop displays window elements in accordance with an ordering termed Z-ordering. Z-ordering identifies the window in the foreground
15 and the various other windows on a tree structure. This tree structure is suitably manipulated when a different window is moved to the foreground, possibly due to a change in focus due to a user clicking on another window and consequently transferring focus. Such transfer of focus can also results in another application or thread getting the input events from the RIT.

20 It is possible for an application to hang, and at the very least give the user the impression of being non-responsive, if an application fails to process its input queue. For instance, ordinarily the main thread of an application manages its windows. Thus, if the main thread of the application is occupied by a task then the entire application may

appear to hang, i.e. be non-responsive. The area on the desktop, i.e., the real estate,
occupied by window elements controlled by its main thread appears to be frozen. A user
may find it difficult to close or manage such an application and may, indeed, lose
significant control over the apparently frozen system. Even a mere delay may result in a
5 drastic action being taken by many a user such as rebooting the system leading to possible
loss of data and even damage to the file structure and the hardware.

SUMMARY OF THE INVENTION

In view of the foregoing, the present invention provides a method and system for
10 executing software application programs that provide a UI for replacing an apparently
frozen application's UI. The replacement UI is responsive enough to permit moving,
sizing, minimization, or closing of the hung window. Furthermore, the frozen or hung
application may be closed by a user without risking rebooting the entire system resulting
in possible data loss or damage to the hardware. This is enabled by implementing,
15 preferably as part of the OS, routines to (a) detect if an application is hung, (b) replace the
UI for the hung application with a ghost UI that permits limited functionality such as
moving, sizing, minimization and closing of the application, and (c) destroying the ghost
UI if the hung application becomes responsive. The destruction of the ghost window is
followed by restoring the application's UI along with updated queue entries.

20 Hung applications are preferably detected by the duration for which they fail to
handle entries queued on their respective input threads. In response to the detection of a
hung application a ghost user interface is created. The ghost UI is preferably provided by
a separate thread dedicated to managing ghost interfaces for hung applications. Creation

of a ghost UI is accompanied by placing a high priority special entry in the queue for the application. This results in detection of the responsiveness of the application by its handling of the high priority special entry ahead of other queue entries. This, in turn, triggers the restoration of the application's UI and destruction of the associated ghost UI if the application becomes responsive again. To the user the effect is of watching a user interface change in some respects while retaining functionality related to managing the window and closing the frozen application using the familiar commands. Furthermore, the facility for caching input for the application makes the transformation smooth with few discontinuities since there are few "lost" commands while the ghost user interface is operative. The extent and nature of caching is implementation specific and affects the smoothness of the transition to a ghost user interface and back again to the application's interface.

As is apparent, the method is generally useful in employing at least one dedicated thread for providing replacement user interfaces for a plurality of threads or applications in response to an appropriate signal and restore the original user interface in response to another signal. Furthermore, the threads may be replaced by independently scheduled code segments that may be executing on different machines or processors without any loss of generality. Of course, the dedicated thread may perform other activities that do not detract from its primary task of providing a replacement interface.

Additional features and advantages of the invention will be made apparent from the following detailed description of illustrative embodiments, which proceeds with reference to the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

While the appended claims set forth the features of the present invention with particularity, the invention, together with its objects and advantages, may best be understood from the following detailed description taken in conjunction with the accompanying drawings of which:

Figure 1 is a block diagram generally illustrating an exemplary computer system on which the present invention resides;

Figure 2 is a schematic providing an overview of a typical window element used to graphically represent an application's real estate in a possible embodiment;

Figure 3 illustrates an application's main window with several child windows elements within it;

Figure 4 illustrates a possible implementation for managing the display of graphic UI elements in a computing system;

Figure 5 illustrates a possible implementation for managing the input in a computing system; and

Figures 6 illustrates high level view of the steps in creating and managing a ghost window;

Figure 7 illustrates a possible implementation for creating a ghost window

Figure 8 illustrates a possible implementation for destroying a ghost window.

DETAILED DESCRIPTION OF THE INVENTION

In the description that follows, the invention will be described with reference to acts and symbolic representations of operations that are performed by one or more computer, unless indicated otherwise. As such, it will be understood that such acts and operations, which are at times referred to as being computer-executed, include the manipulation by the processing unit of the computer of electrical signals representing data in a structured form. This manipulation transforms the data or maintains it at locations in the memory system of the computer, which reconfigures or otherwise alters the operation of the computer in a manner understood by those possessing ordinary skill in the art. The data structures where data is maintained are physical locations of the memory that have particular properties defined by the format of the data. However, while the invention is being described in the foregoing context, it is not meant to be limiting as those of skill in the art will appreciate that various of the acts and operation described hereinafter may also be implemented in hardware.

Turning to the drawings, wherein like reference numerals refer to like elements, the invention is illustrated as being implemented in a suitable computing environment. Although not required, the invention will be described in the general context of computer-executable instructions, such as program modules, being executed by a personal computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, multi-processor systems, microprocessor based or programmable consumer electronics, network

PCs, minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

With reference to Fig. 1, an exemplary system for implementing the invention includes a general purpose computing device in the form of a conventional personal computer 20, including a processing unit 21, a system memory 22, and a system bus 23 that couples various system components including the system memory to the processing unit 21. The system bus 23 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. The system memory includes read only memory (ROM) 24 and random access memory (RAM) 25. A basic input/output system (BIOS) 26, containing the basic routines that help to transfer information between elements within the personal computer 20, such as during start-up, is stored in ROM 24. The personal computer 20 may further include a hard disk drive 27 for reading from and writing to a hard disk 28, a magnetic disk drive 29 for reading from or writing to a removable magnetic disk 30, and an optical disk drive 31 for reading from or writing to a removable optical disk 32 such as a CD ROM or other optical media.

The hard disk drive 27, magnetic disk drive 29, and optical disk drive 31 are connected to the system bus 23 by a hard disk drive interface 33, a magnetic disk drive interface 34, and an optical disk drive interface 35, respectively. The drives and their associated computer-readable media provide nonvolatile storage of computer readable

instructions, data structures, program modules and other data for the personal computer

20. Although the exemplary environment described herein employs a hard disk 28, a removable magnetic disk 30, and a removable optical disk 32, it will be appreciated by those skilled in the art that other types of computer readable media which can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks, Bernoulli cartridges, random access memories, read only memories, and the like may also be used in the exemplary operating environment.

A number of program modules may be stored on the hard disk 28, magnetic disk 30, optical disk 32, ROM 24 or RAM 25, including an OS 36, one or more applications programs 37, other program modules 38, and program data 39. A user may enter commands and information into the personal computer 20 through input devices such as a keyboard 40 and a pointing device 41. Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 21 through a serial port interface 42 that is coupled to the system bus. Increasingly, such devices are being connected by the next generation of interfaces, such as a universal serial bus (USB) 43 with a root hub/Host 44, and to which other hubs and devices may be connected. Other interfaces that may be used include parallel ports, game ports, and the IEEE 1394 specification available at

<http://www.standards.ieee.org/catalog/bus.html#1394-1995>. A monitor 45 or other type of display device is also connected to the system bus 23 via an interface, such as a video adapter 46. In addition to the monitor, personal computers typically include other peripheral output devices.

The USB connections illustrate its utility. A keyboard 47, a pointing device 48 and another hub, hub-1 49, are connected to the root hub/Host 44. Hub-1 49 is further connected to another hub, hub-2, 50, scanner 51, monitor 52, camera-1 53, and modem 54. It should be understood that additional cameras and devices may be directly
5 connected to the computer instead of a USB. Thus, the system depicted is capable of communicating with a network and sending/receiving audio, video and data.

The personal computer 20 may operate in a networked environment using logical connections to one or more remote computers. The types of connections between networked devices include dial up modems, e.g. modem 51 may be directly used to
10 connect to another modem, ISDN, xDSL, cable modems, wireless and include connections spanning users connected to the Internet. The remote computer may be another personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the personal computer 20 in Fig. 1. The logical connections depicted in
15 Fig. 1 include a local area network LAN 55 and a wide area network WAN 56. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the personal computer 20 is connected to the local network 51 through a network interface or adapter 53. When used
20 in a WAN networking environment, the personal computer 20 typically includes a modem 54 or other means for establishing communications over the WAN 52. The modem 54, which may be internal or external, is connected to the system bus 23 via the serial port interface 46. In a networked environment, program modules depicted relative

to the personal computer 20, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used. In particular, distributed computing based on dynamic networks that can reconfigure themselves with a device providing functionality, such as a video display, to another device is intended to be included.

In the description that follows, the invention will be described with reference to acts and symbolic representations of operations that are performed by one or more computers, unless indicated otherwise. As such, it will be understood that such acts and operations, which are at times referred to as being computer-executed, include the manipulation by the processing unit of the computer of electrical signals representing data in a structured form. This manipulation transforms the data or maintains it at locations in the memory system of the computer, which reconfigures or otherwise alters the operation of the computer in a manner well understood by those skilled in the art. The data structures where data is maintained are physical locations of the memory that have particular properties defined by the format of the data. However, while the invention is being described in the foregoing context, it is not meant to be limiting as those of skill in the art will appreciate that various of the acts and operation described hereinafter may also be implemented in hardware.

Table 1 presents a sampling of graphical elements or windows used as UI elements in addition to the familiar window interface. This list is not intended to be exhaustive, and instead, is only illustrative of the some of the behaviors made possible by different graphical UI elements. Different OSs and applications include additional icons,

- bitmaps and graphical shapes that also function as UI elements. Often the UI elements facilitate in the generation of events that are subsequently handled by the application or the OS or even plug-in routines. It is to be noted that additional graphical controls, including those with audio and video properties are being developed and being
- 5 continually released. Thus, “graphical controls” and other similar terms should be understood to include audio/video capable elements.

Description	Options	Notifications
<u>Status bar</u> : displays information defined by an application.	Simple-mode – one section; and Multi-mode – displaying more than one type of information in their respective sections.	May generate events corresponding to mouseovers or mouse clicks over a particular section.
<u>Button</u> : a bitmap displaying selected text.	Can associate text or selected bitmaps to modify its appearance.	May generate events corresponding to mouseovers or mouse clicks over a particular section.
<u>Message box</u> : displays a text message.	Can associate selected text.	May generate events corresponding to mouseovers or mouse clicks over a particular section.
<u>Tool bar</u> : displays a collection of buttons in a bar.	Can associate selected text or pictures.	May generate events corresponding to mouseovers or mouse clicks over a particular section.

Table 1

Description	Options	Notifications
<u>Tool tip</u> : displays text explaining the tool function.	Can associate selected text.	May generate events corresponding to mouseovers over a particular button or another graphic.
<u>Trackbar</u> : displays a scrolling control with a slider and a set of notches.	Can associate text and numbers with the scale.	May generate events corresponding to mouse-dragging events.
<u>Spin box</u> : displays arrows in an updown control.	Can specify location and size.	May generate events corresponding to mouse-clicks over a particular section.

Table 1 (continued)

Figure 2 is an illustration of a customary window element 60, also termed a “top-level window,” used to represent the UI for an application. The window element 60 has a title and logo 62 in the system area along with system buttons for minimization 64, restore 66 and close 68 functions. The window element encloses a client area 70 that is available for placing text and/or graphics. It should be noted that the content of the title and logo 62 is also controllable by the application. Furthermore, the window element 60 has borders 72 that can usually be manipulated to change the real estate occupied by the window element 60 on the desktop. Similarly, the corner 74 can be used to change the real estate by changing the base and the side of the window element 60 by dragging it with a pointing device.

Figure 3 is an illustration of a top-level window 80 with several child windows contained therein. These child windows may typically be controlled by, i.e., created by, the application thread that created the parent top-level window or be controlled by another thread that acts in concert with the thread controlling the top-level window. The top level window 80 has a left system button 82 for opening a drop down menu to select familiar options such as "move," "close" etc. The top-level window also has a title bar 84 and system buttons. The system buttons are represented by familiar standard shapes illustrated by the "minimize" button 86, "restore" button 88, and the "close" button 90. The child windows contained in the client area of the top-level window 80 include other windows 92 and a dialog box 94 to illustrate the variety of graphical symbols that may be present. Some of the child window elements contain the familiar system buttons.

In order to understand the relationship between the displayed window elements, and the organization of the computing environment, it is useful to understand the thread-based computing environment discussed previously. In graphical user interfaces (GUIs) there are many different tasks that generate information to be displayed to the user. Similarly there is a constant stream of inputs from a variety of sources such as the keyboard, the pointing devices and the like. All of these need to be organized in order to allow programs to receive input and send output in a consistent manner without losing or misinterpreting data.

Figure 4 illustrates an embodiment of a system for managing a display, which is usually a desktop, with a plurality of window elements corresponding to applications and system components. Advantageously, a desktop thread is used to manage the desktop. The window elements on a desktop are managed by the window manager, which arranges

them in a tree structure termed Z-ordering. The window in the foreground is the root of the tree and the remaining windows are occluded or covered to different extents by other windows depending on their position in the tree. The foreground window is also the window in focus. The selection of another window to be in focus results in manipulation
5 of the Z-order to reflect the change and an updating of the desktop.

In one of the embodiments schematically described in figure 4, if the user shifts the focus to another window, the system adjusts the Z-order tree accordingly and generates a signal to the relevant applications to request a repaint() operation so that the changes to the desktop can be implemented. Such requests can advantageously placed in
10 the virtual queue for a given application. This arrangement results in the applications tracking the details of the graphical output.

Step 100 in figure 4 shows a change in the Z-ordering made by one of the system threads in concert with the window manager, which tracks the different windows. The window manager identifies the windows affected by the change (step 102) along with the
15 threads or applications that own the windows. In order to update the desktop, a request to repaint() is made to the owners of the various window elements (step 104). This request is preferably tagged with the time at which the request is made. The repaint request is handled by providing the necessary graphical information by calling relevant APIs (step 106). The application or thread then calls repaint() and thus supplies the result to the
20 system by calling graphics APIs (step 110) while noting the time at which the request to repaint was handled (step 108). These time stamps allow an estimate to be made of the promptness with which the application handles the requests queued in its message or event queue for handling. Finally, particularly in systems using " α " in addition to the

customary colors, an appropriate system thread composites the final image (step 112). It should be noted that the last step need not be found in many implementations, and instead the application may supply the repainted input with no need to composite it with other images.

5 In some implementations, layered windows are implemented so that the repainting operation may include compositing. In a nutshell, the layered window feature is implemented by maintaining an image of the an area of the screen occupied by the layered window that contains all the bits not belonging to the layered window. When either a layered window repaints or there are drawing operations in the specified area of the
10 screen, the recompositing operation occurs in the context of the of the thread causing the changes. In any event, the precise details for implementing compositing to allow use of one or more transparency parameters, such as α , in course of repainting are implementation details particular to some of the many embodiments enabled by the invention described here. Such details are intended to be within the scope of the claimed
15 invention.

The input to the computing system is handled by the RIT, as is illustrated in figure 5. RIT receives (step 120) and determines which application, if any, is the intended target for the input (step 122). It accordingly forwards the input to the relevant applications (step 124) by placing it in their respective virtual input queue. A helpful rule is to send
20 the input to the application in focus for receiving input. This arrangement dissociates the processing of the input from the actual receipt of the input by implementing a virtual input queue at the application and removes the potential for system-wide bottlenecks due to a single task's failure to handle the input. Preferably, in an embodiment, each entry in

this queue is marked with a time stamp indicating the time when it was placed in the virtual input queue for a particular application by RIT. The queue also has a time stamp to indicate when a queued entry was last handled by the application (step 126).

5 Suppose that an application has stopped processing messages in its input queue while waiting for a network call or some computationally expensive operation to complete. This scenario may make the system appear to be hung since all of the input is sent to the application's virtual input queue and elicits no response from the busy application while the foreground window makes a shift in the focus difficult by obscuring other windows behind it.

10 Furthermore, even if another window is selected to be in the foreground, it is difficult to regenerate the desktop if an application window occupying a significant of the desktop real estate is non-responsive to requests to repaint(). This may result in an imperfectly painted desktop that looks strange and appears to be, at least partially, non-responsive.

15 In light of these considerations, Figure 6 illustrates a high level illustration of a possible embodiment in accordance with the invention for making the system appear more responsive, and in particular making the user interface responsive enough to generate a desktop that is not frozen. This task is advantageously accomplished using a three-step procedure. The first step is detection of a hung application (step 130). There
20 are many methods that are suitable for determining if an application is hung. While the user notices a non-responsive system, the inability of an application to handle input queued in its virtual input queue is often the real reason behind the non-responsive user interface. A system thread, e.g., a desktop thread or RIT, can detect a hung application by

noting if the queued entries are being handled within a set threshold period of time.

Alternatively, different types of entries can have their own thresholds or the number of entries in the queue and other factors like the application or thread priority could also be considered in deciding if an application was actually hung. Different methods of

5 determining whether an application is hung, including the user having the option of inputting a suitable signal, are intended to be included in step 130.

If an application is hung then the system thread determines if a ghost thread exists (step 132). The desktop thread, for instance, may make such a determination if a request to repaint is queued for too long. Similarly, the RIT or even another thread may be
10 charged to make such a determination. If there is no ghost thread, then a ghost thread is created (step 134). It should be noted that although in this embodiment a ghost thread is used to control the creation and destruction of all ghost windows, this is not intended to be a limitation. The task of managing ghost windows could be handled by another thread or more than one thread in different embodiments in accordance with the invention. If a
15 ghost thread exists then it is asked to create a ghost window (step 136) using the location and properties of the frozen window belonging to the hung application. These properties are advantageously obtained from the system threads managing the display.

The ghost window preferably covers the frozen window belonging to the hung application while the hung window itself is hidden from the end-user's view. Now, the
20 desktop thread does not need to use the input from the hung application to update the display and instead uses the ghost window by directing a repaint request to the ghost thread instead. The ghost thread calls the APIs to create the ghost window and then calls repaint() (step 138). This permits a complete update of the display (step 140) and

handling of the input that otherwise would have to be handled by the now hung application. Similarly, mouse movements and mouse-clicks corresponding to the area occupied by the frozen window are handled by the ghost thread. In part, the system buttons come back to life at least as far as the user is concerned. The ghost window, corresponding to the previously frozen window, can be minimized, maximized, resized, moved and even closed. Each of these actions has different consequences to complete the simulation of the application window element. The close system button preferably terminates, i.e., results in the hung application exiting although other actions are possible. Additional or fewer functions may be implemented in the ghost UI. In effect, the ghost UI is a place saver for the application's UI.

In an embodiment in accordance with the invention, the events and messages for the ghost window are preferably divided into two sets depending on whether they should be handled by the ghost thread in lieu of the application or cached for later handling by the application. This is advantageously implemented by treating input with the ghost window in focus as being intended for the application except for some input events that may be handled by the ghost thread. The remaining input events are cached for subsequent forwarding to the hung application. Furthermore, inputs such as moving the ghost window are also forwarded to the application to restore the application window in the position occupied by the ghost window if and when the application ceases to be non-responsive.

A system thread detects if the hung application starts handling its queued entries (step 142). Advantageously the renewed responsiveness of an application may be detected without requiring extensive monitoring by placing a high priority special entry in

the queue for the application. Thus, when the hung application completes its task and is able to handle queue entries, it first responds to the high priority entries. The high priority special entry merely requires it to send a message, directly or indirectly, to a thread, e.g., the ghost thread in order to initiate replacement of the place saver ghost window by the real application window. The ghost thread detects that the application is no longer non-responsive, i.e., hung, (step 142) and destroys the ghost window, initiates restoration of the application window by placing appropriate requests to repaint and forwarding cached input to the application queue (step 144).

Figure 7 illustrates an implementation of generating the ghost window. A system thread receives or generates events/messages for handling by an application (step 150). The system thread time stamps the event/message with the time at which it is placed in the queue for the application (step 152). The system thread checks the timestamp for the queue to determine the time when the last queue entry was handled by the application. If the application has not handled queue entries for greater than a threshold time (step 154) then it is treated as being hung (step 156). Else, the application is continues (step 158).

If the application is hung then it is determined whether a ghost thread exists (step 160). If there is no ghost thread then one is created by a system management thread (step 162). Else, the ghost thread is requested to generate another instance of a window to replace the window corresponding to the hung application (step 164) followed by placing a high priority special entry in the application's queue (step 166). In this scheme the system thread could be the desktop thread, RIT or another thread. Furthermore, the system managing thread could be the desktop thread as well in some embodiments.

Figure 8 further illustrates an implementation in accordance with the invention for restoring a newly responsive application's window. At step 170, a high priority special entry is placed, and naturally moved ahead of other entries in a hung application's queue.

When the hung application gets over the bottleneck that made it non-responsive, it starts
5 handling entries in its queue at step 172. It first handles the high priority special entry, and consequently sends a message, directly or indirectly, to a ghost thread. The ghost thread destroys the ghost window corresponding to the now responsive application (step 176) along with updating cached entries for the application to remove redundant entries (step 178). Such redundant entries may include intermediate move commands executed
10 on the ghost window and now are replaced by the last location of the ghost window as the place to restore the application window. A suitable last state may even be as a minimized window element. The updated cached entries are added to the queue for the application (step 180). This has the effect of enhancing continuity between the ghost window and the application window. The user is only aware of the application window changing while
15 the application was busy due to the absence of detailed application specific information in the client area of the frozen application window. In some embodiments the ghost window may even identify itself as a ghost window with limited functionality. But the application does not become non-responsive as far as the user is concerned because it can apparently be minimized out of the way or even closed to smoothly terminate it.

20 Since there is no reason to keep the ghost thread if there are no ghost windows, at least in many embodiments, a check is made to determine if there are any more ghost windows remaining (step 182). If there are no more ghost windows then the ghost thread exits, else it continues to manage its other ghost windows.

Figure 9 illustrates the operation of the close command in an embodiment of the invention. The ghost window receives a close command at step 190. The ghost thread responsively initiates a forced closure of the hung application at step 192. Such a forced termination may be carried out, for instance, by invoking the taskman.exe in the
5 "WINDOWS®" OS environment manufactured by the "MICROSOFT®" Corporation of Redmond, Washington. Since there is no need for forwarding entries to the hung application, the cached entries are discarded (step 194) and the ghost window is destroyed (step 198). If there are no other ghost windows being managed by this ghost thread (step 200) then the ghost thread exits (step 202).

10 Figure 10 illustrates a possible system in accordance with the invention. In a computing system with ghost interfaces 210 an application 212 utilizes System Services 214 for providing a user interface 216. The system services may include a desktop thread and the graphical APIs used to generate a user interface. Application 212 has at least one queue 218 for receiving input from users and other programs, messages as part of a client-
15 server architecture and the like. The queue 218 has two way communications with code 220 for detecting if application 212 is hung. Code 220 may be part of a system thread, e.g., the desktop thread or RIT or be a separate component. Code 220 triggers the ghost thread 222 resulting in creation of a ghost user interface 224 which can communicate with the System Services 214 for acquiring data to replace the Application's user
20 interface 216. System Services 214 also redirects input and focus along with providing functionality for executing system buttons such as move and close commands in lieu of the Application. Ghost thread 222 or code 220 can place a high priority special entry 226 in the queue 218. Handling of this entry by a hung application indicates that the

application is responsive and results a message being sent to the ghost thread 222 and the instructions 228 for destroying the ghost user interface, which, in turn initiate the destruction of the ghost user interface 224 by the ghost thread. Instructions 228 can advantageously be part of the ghost thread in some embodiments. It should be understood that not all possible system components are described in Figure 10 in the interest of clarity.

It should be noted that the procedure described is quite general and can be advantageously implemented in contexts that do not require an application to become non-responsive. Thus, a first window element, controlled by a scheduled code segment, e.g., a thread, providing replacement window elements, may be invoked to replace a second window element in response to a flip-window signal. Another flop-window signal can result in replacing the first window with the second window. Events associated with the windows can be handled analogously to that described herein for hung applications. In the case of embodiments replacing hung windows, the flip-windows signal corresponds to detection of a hung application or thread and the flop-window signal, then, corresponds to detection of the hung application or the hung-thread becoming responsive again. Furthermore, in a networked computing environment it may be possible to use applications or code executing on different processors or even machines to provide the independently scheduled functionality of threads to provide ghost or replacement user interface elements for a variety of applications. Thus, a scheduled code segment is a more general concept than the more familiar thread.

While much of the description of embodiments in accordance with the invention are in the "WINDOWS®" OS environment, this description is not intended to exclude

other OSs such as the "MACINTOSH®," "SOLARIS®" and other UNIX based platforms along with distributed computing.

All of the references cited herein, including patents, patent applications, and publications, are hereby incorporated in their entireties by reference.

5 In view of the many possible embodiments to which the principles of this invention may be applied, it should be recognized that the embodiment described herein with respect to the drawing figures is meant to be illustrative only and should not be taken as limiting the scope of invention. For example, those of skill in the art will recognize that the elements of the illustrated embodiment shown in software may be implemented
10 in hardware and vice versa or that the illustrated embodiment can be modified in arrangement and detail without departing from the spirit of the invention. Therefore, the invention as described herein contemplates all such embodiments as may come within the scope of the following claims and equivalents thereof.

CLAIMS

We claim:

1. A method for managing a user interface in a multithreaded computing environment, the user interface comprising a plurality of user interface elements wherein
5 a first user interface element in the plurality of user interface elements corresponds to a first application, wherein furthermore the first application having a first thread having control over the first user interface element, method comprising the steps of:
signaling a hung state for the first application; (if there is a delay greater than a
predetermined threshold in handling events queued in an event queue for handling by the
10 first application);
creating a ghost user interface element, with a ghost thread, responsively to a hung signal indicating the first application is in the hung state wherein the ghost user interface element replaces the first user interface element in the user interface;
placing a high priority special event in the event queue for handling by the first
15 application wherein handling of the special event generates a activity-detect message; and
detecting the wakeup message and responsively to the activity-detect message replacing the ghost user interface element by the first user interface element.
2. The method of claim 1 wherein the step of signaling includes signaling a hung
20 state if the first application does not handle an event in an event queue for at least a predetermined duration.
3. The method of claim 1 wherein the step of detecting the activity-detect message further comprises releasing, responsively to the activity-detect message, resources used by
25 the ghost user interface element.
4. The method of claim 1 wherein the step of creating the ghost user interface element includes creating the ghost thread, which, in turn, creates the ghost user interface element.

30

5. The method of claim 4 wherein the step of creating the ghost thread includes determining if a ghost thread exists and creating a ghost thread if the ghost thread does not exist.

5 6. The method of claim 1 wherein the step of creating the ghost user interface element includes creating the ghost user interface element in the area occupied by the first user interface element.

10 7. The method of claim 6, the method further having the step of directing input events corresponding to the area covered by the ghost user interface element, which includes at least some of the area covered by the first user interface element, to the ghost thread.

15 8. The method of claim 7, the method further having the step of having the ghost thread handle at least one event corresponding to the area covered by the ghost user interface.

20 9. The method of claim 7 having the step of having the ghost thread cache at least one event corresponding to the area covered by the ghost user interface to the ghost thread for handling by the first application.

25 10. The method of claim 7, the method further having the step of having the ghost thread handle at least one event corresponding to a minimization operation in the area covered by the ghost user interface element.

11. The method of claim 7, the method further having the step of having the ghost thread handle at least one event corresponding to a resizing operation in the area covered by the ghost user interface element.

12. The method of claim 7, the method further having the step of having the ghost thread handle at least one event corresponding to a close operation in the area covered by the ghost user interface element.

5 13. The method of claim 7, the method further having the step of having the ghost thread handle at least one event corresponding to a move operation in the area covered by the ghost user interface element.

10 14. The method of claim 7, the method further having the step of having the ghost thread cache at least one event corresponding to a keyboard input corresponding to the ghost user interface element.

15 15. The method of claim 7, the method further having the step of having the ghost thread handle at least one event corresponding to a keyboard input corresponding to the ghost user interface element.

16. A method for replacing a first user interface element created by a first application by a first substitute user interface element created by a scheduled code path in a multithreaded computing environment, the method comprising the steps of:
20 detecting a first flip-window signal while the first user interface element is being displayed;
creating, responsively to the first flip-window signal, the first substitute user interface element in the context of the scheduled code path;
superimposing the first substitute user interface element over a first-user-interface-
25 element-occupied real estate; and
caching input corresponding to the first substitute user interface element wherein the cached input is subsequently handled by the first application.

30 17. The method of claim 16 wherein the step of creating includes creating the scheduled code path if the scheduled code path does not exist.

18. The method of claim 16, the method further having the step of replacing a second user interface element created by a second application by a second substitute user interface element created by the scheduled code path responsively to a second flip-window signal.

19. The method of claim 16, the method further having the step of painting over the first user interface element responsively to the first flip-window signal.

20. The method of claim 16 wherein the step of forwarding includes input comprising a plurality of events and/or a plurality of messages.

21. The method of claim 17, the method further having the step of caching input corresponding to the application for forwarding to the application in response to a first flop-window signal.

22. The method of claim 18, the method further having the steps of selecting and discarding at least some of the cached input.

23. The method of claim 19 wherein the step of forwarding includes forwarding the input to the scheduled code path.

24. The method of claim 20, the method further having the step of handling at least some of the input by the scheduled code path.

25. The method of claim 24, the method further having the step of the scheduled code path handling at least one event in the input by terminating the first application.

26. The method of claim 16, the method further having the step of placing a special message/event in a queue for the first application to generate a flop-window signal.

27. The method of claim 16 wherein the scheduled code path is a thread.

28. A method of using a designated scheduled code path for providing substitute user
5 interfaces for replacing user interfaces corresponding to a plurality of scheduled code
paths, the comprising the steps of:

generating a first flip-window signal corresponding to a first scheduled code path;

generating a second flip-window signal corresponding to a second scheduled code path;

10 replacing, responsively to the first flip-window signal, a first window controlled by the
first scheduled path with a first substitute window controlled by the designated scheduled
code path; and

replacing, responsively to the second flip-window signal, a second window controlled by
the second scheduled path with a second substitute window controlled by the designated
scheduled code path.

15 29. The method of claim 28, the method further having the steps of generating a first
flip-window signal; and replacing, responsively to the first flip-window signal, a first
substitute window controlled by the designated scheduled code path with the first window
controlled by the first scheduled path.

20 30. A multithreaded computing system for executing a first application having a first
user interface, wherein if the first application is non-responsive to user input, the first user
interface is replaced by a first ghost user interface, the system comprising:

a non-responsive-application detecting code for detecting a non-responsive application;

25 a ghost thread for creating the first ghost user interface to replace the first user interface
responsively to detection of a non-responsive application;

a high priority special entry in a queue for the first application for detecting if the first
application is responsive; and

30 a plurality of computer executable instructions for destroying the first ghost user interface
responsively to handling of the high priority special entry by the first application.

31. The system of claim 30 further having a cache of events and messages for handling by the first application wherein the events and messages are directed at the ghost thread.

5

32. The system of claim 30 further having a second non-responsive application with a second user interface and a second ghost user interface created by the ghost thread.

33. The system of claim 30 further having a responsive-application detecting code for detecting when a non-responsive application becomes responsive.

10

34. The system of claim 33 further having a responsive-application user interface restoring code for replacing the first ghost user interface with the first user interface responsively to detecting that the first application has become responsive.

15

35. A computer readable media having computer executable instructions for carrying out the steps of a method for managing a user interface in a multithreaded computing environment, the user interface comprising a plurality of user interface elements wherein a first user interface element in the plurality of user interface elements corresponds to a first application, wherein furthermore the first application having a first thread having control over the first user interface element, method comprising the steps of: signaling a hung state for the first application; (if there is a delay greater than a predetermined threshold in handling events queued in an event queue for handling by the first application); creating a ghost user interface element, with a ghost thread, responsively to a hung signal indicating the first application is in the hung state wherein the ghost user interface element replaces the first user interface element in the user interface; placing a high priority special event in the event queue for handling by the first application wherein handling of the special event generates a activity-detect message; and

20

25

detecting the wakeup message and responsively to the activity-detect message replacing the ghost user interface element by the first user interface element.

ABSTRACT OF THE INVENTION

A method and system are provided for managing application programs with non-responsive user interfaces, possibly due to a bottleneck in handling events/messages. The method and system make an apparently frozen application user interface responsive so as to permit alternative tasks be executed or close the hung application. This is enabled by implementing routines to (a) detect if an application is hung, (b) replace the user interface for the hung application with a ghost interface, implemented by a separate thread, that permits system functionality such as sizing, minimization and closing of the application, and (c) destroying the ghost interface if the hung application becomes responsive again along with restoring the application's user interface along with updated event queues. Furthermore, creation of the ghost user interface is accompanied by placing a high priority special event in the hung application's queue to detect renewed responsiveness without requiring explicit monitoring.

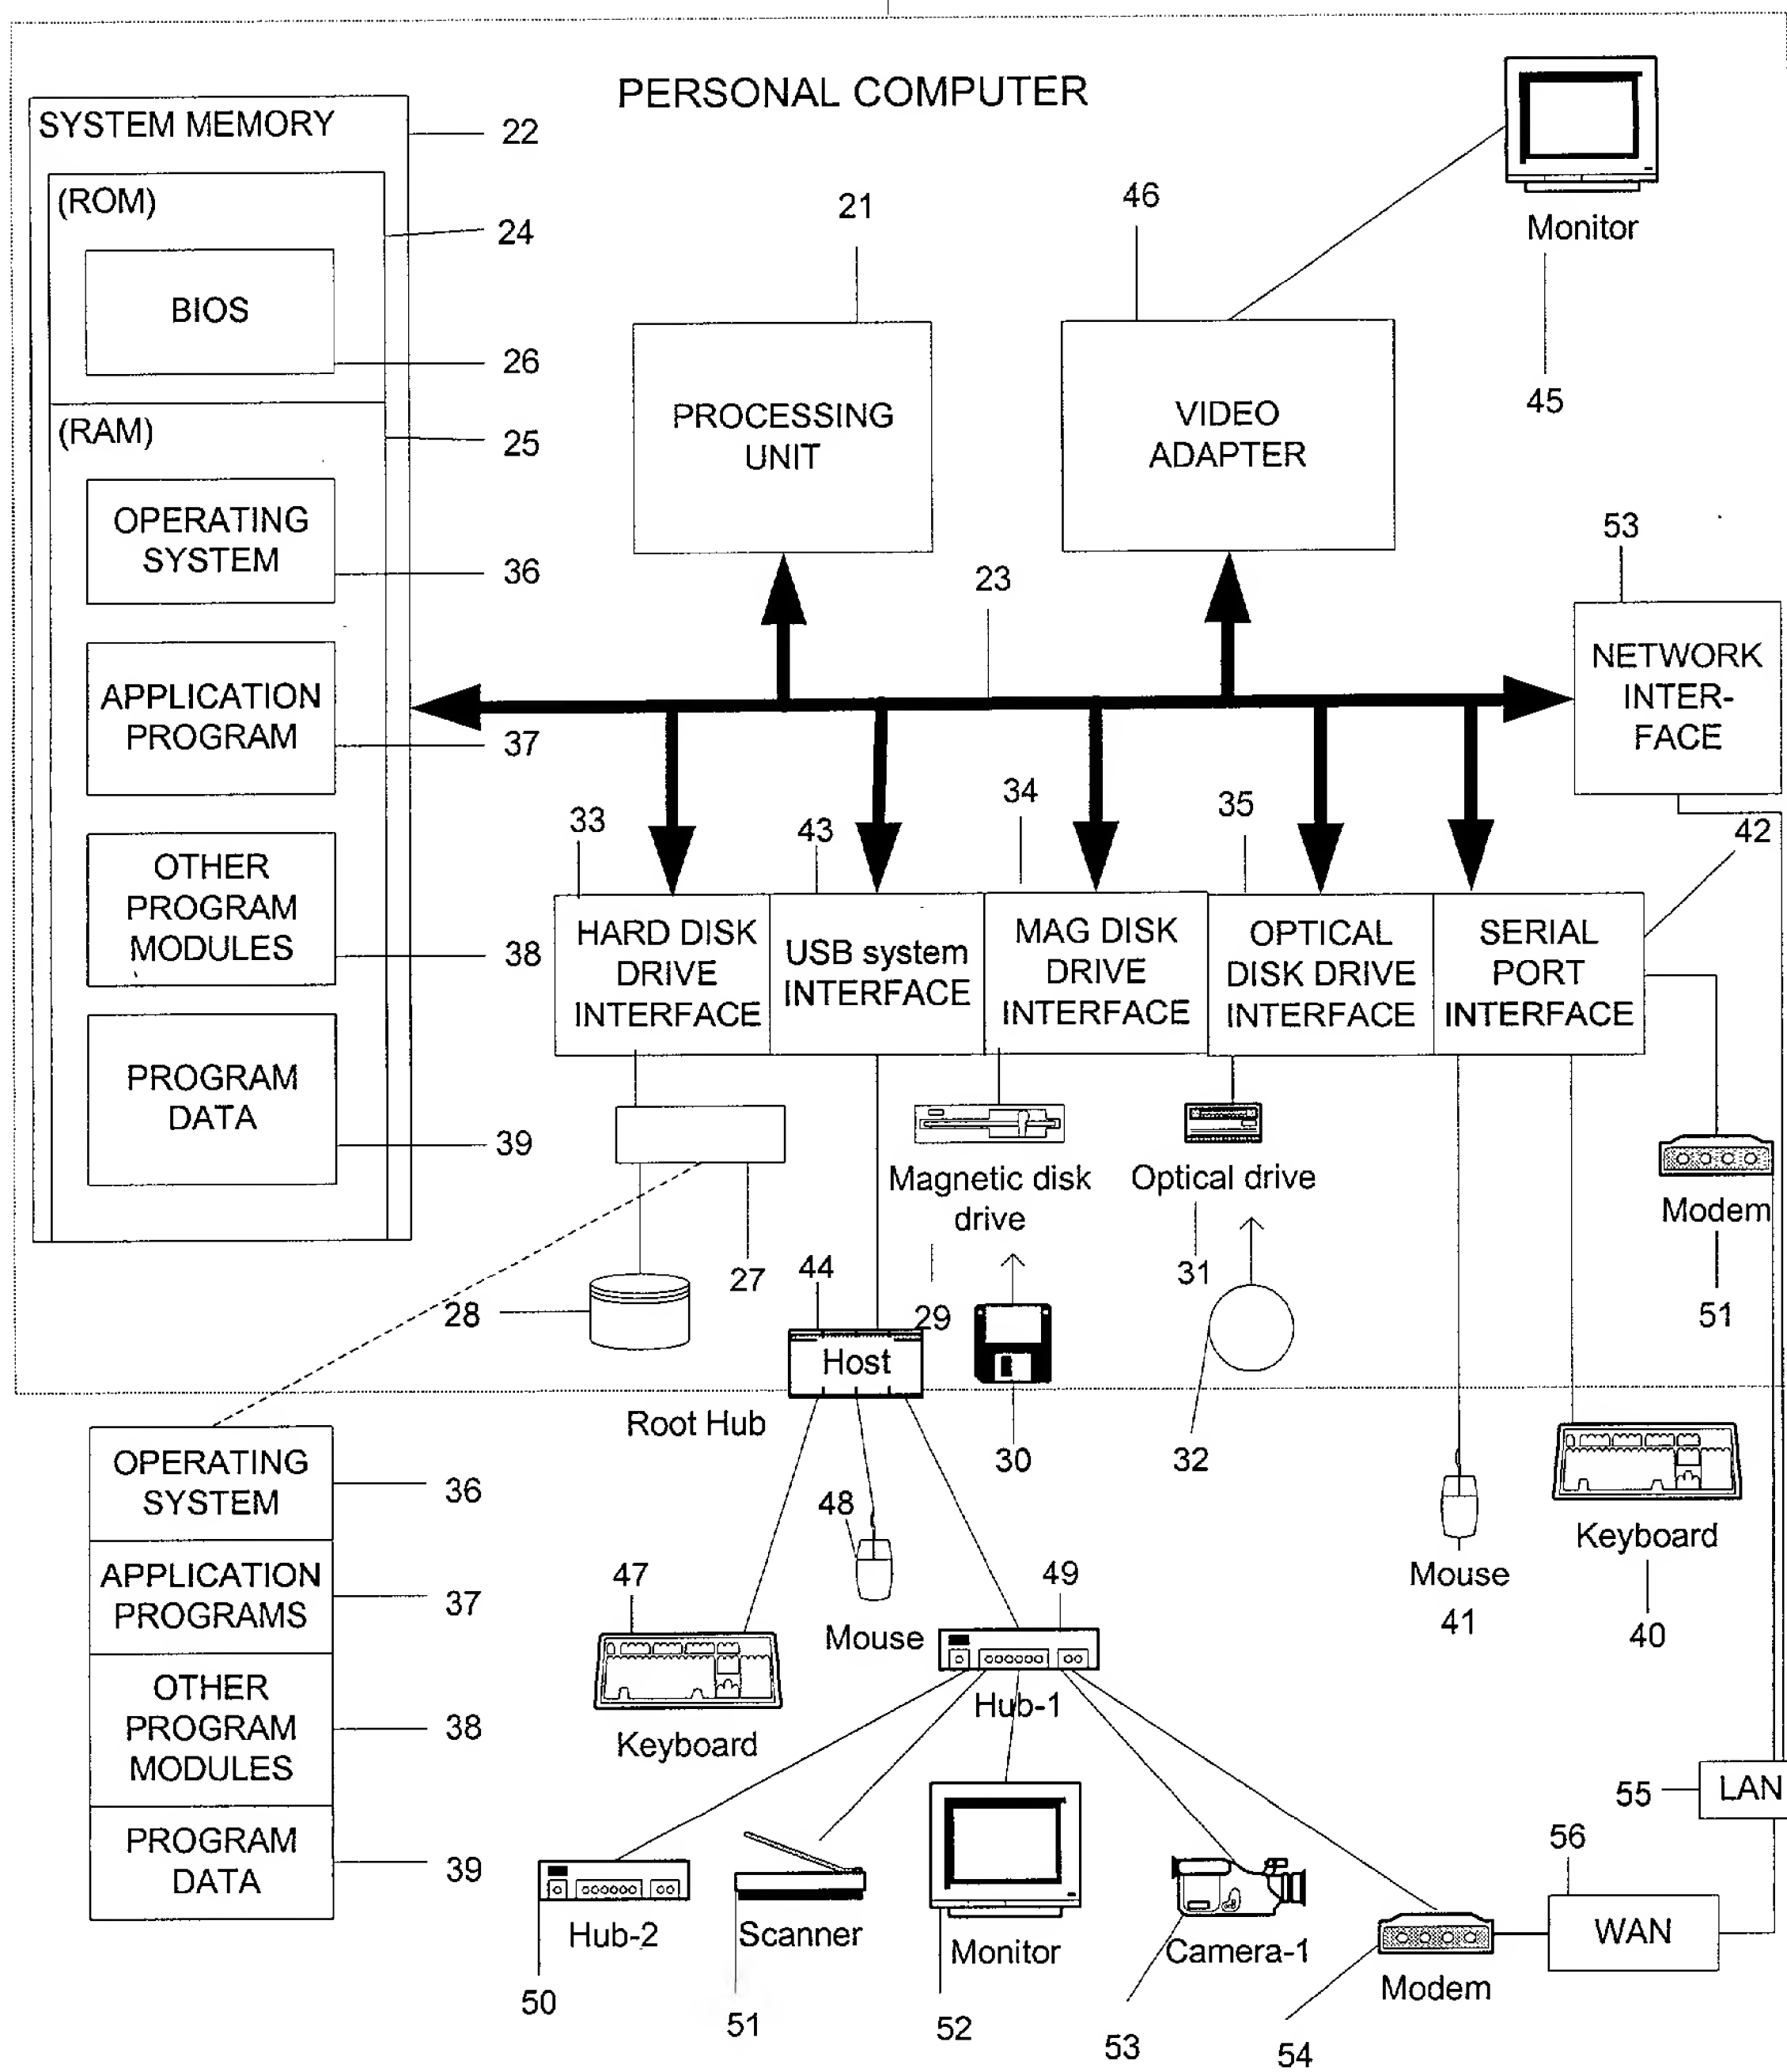


Figure 1




Figure 2




FIGURE 3

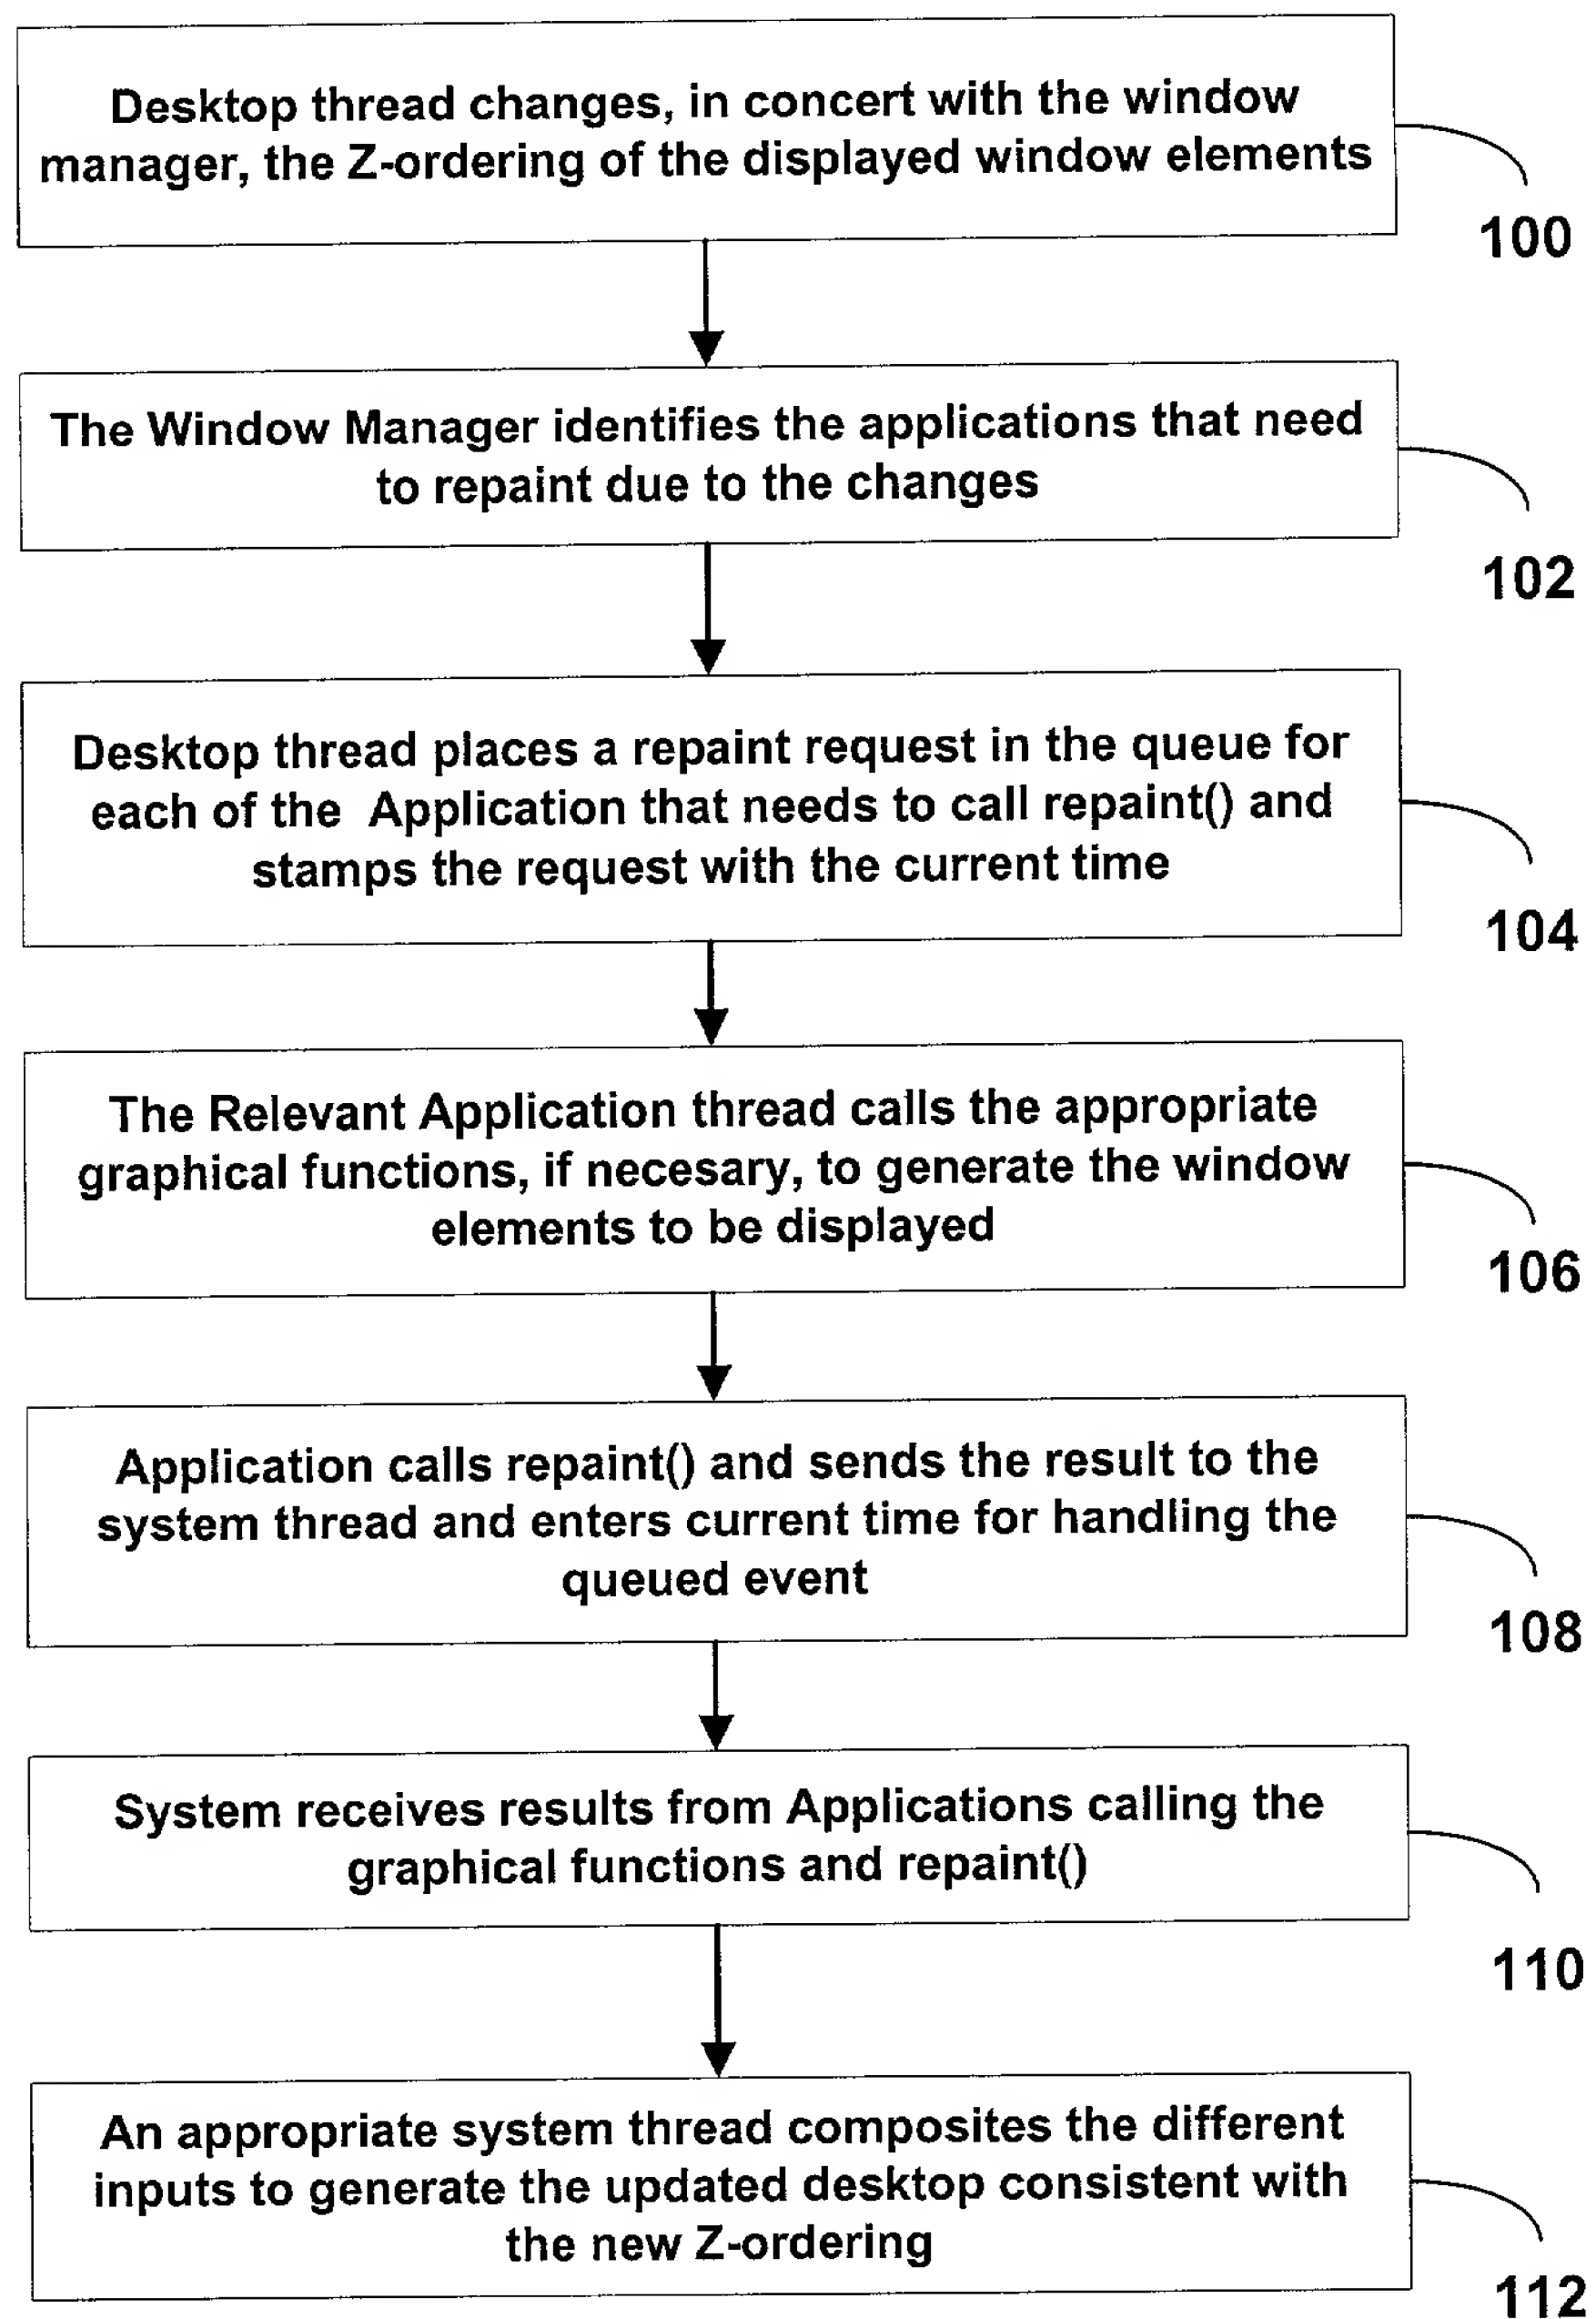


FIGURE 4

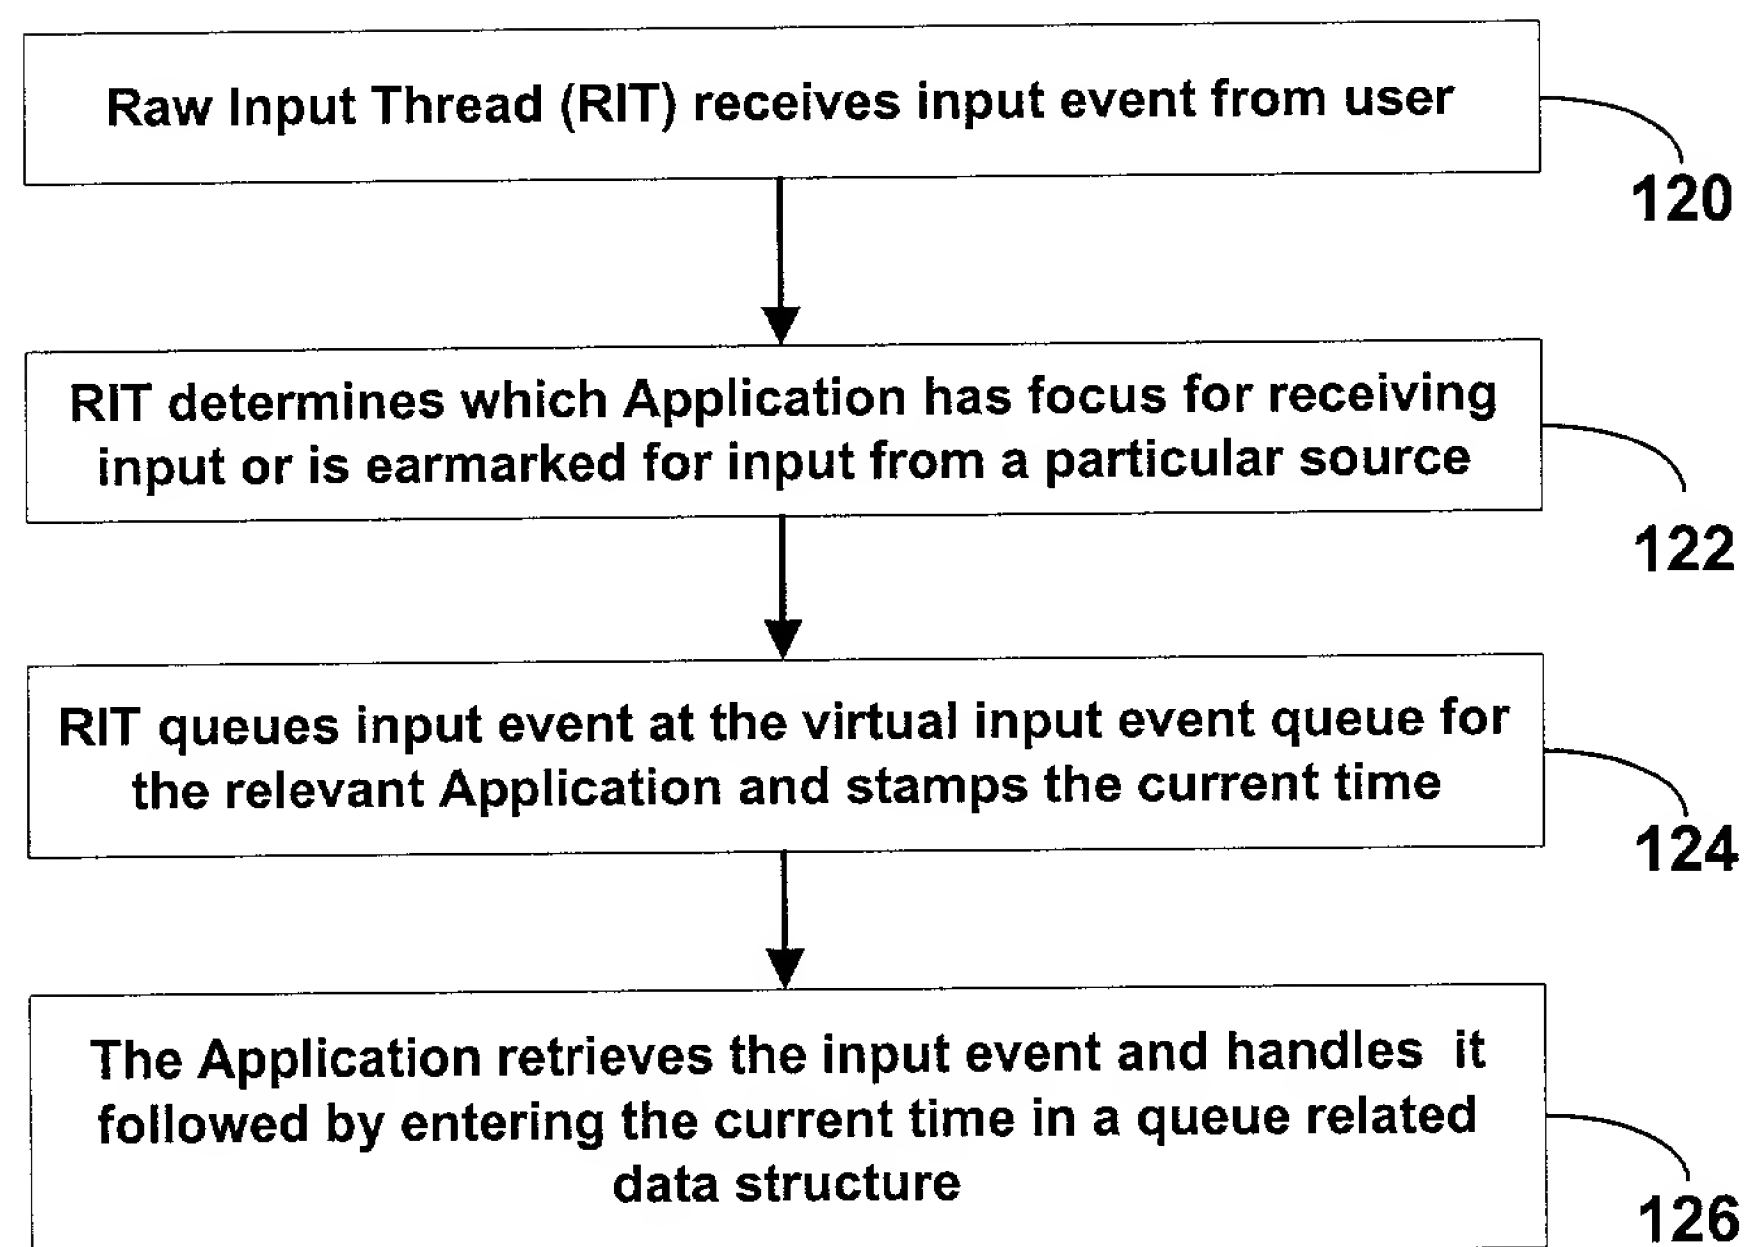


FIGURE 5

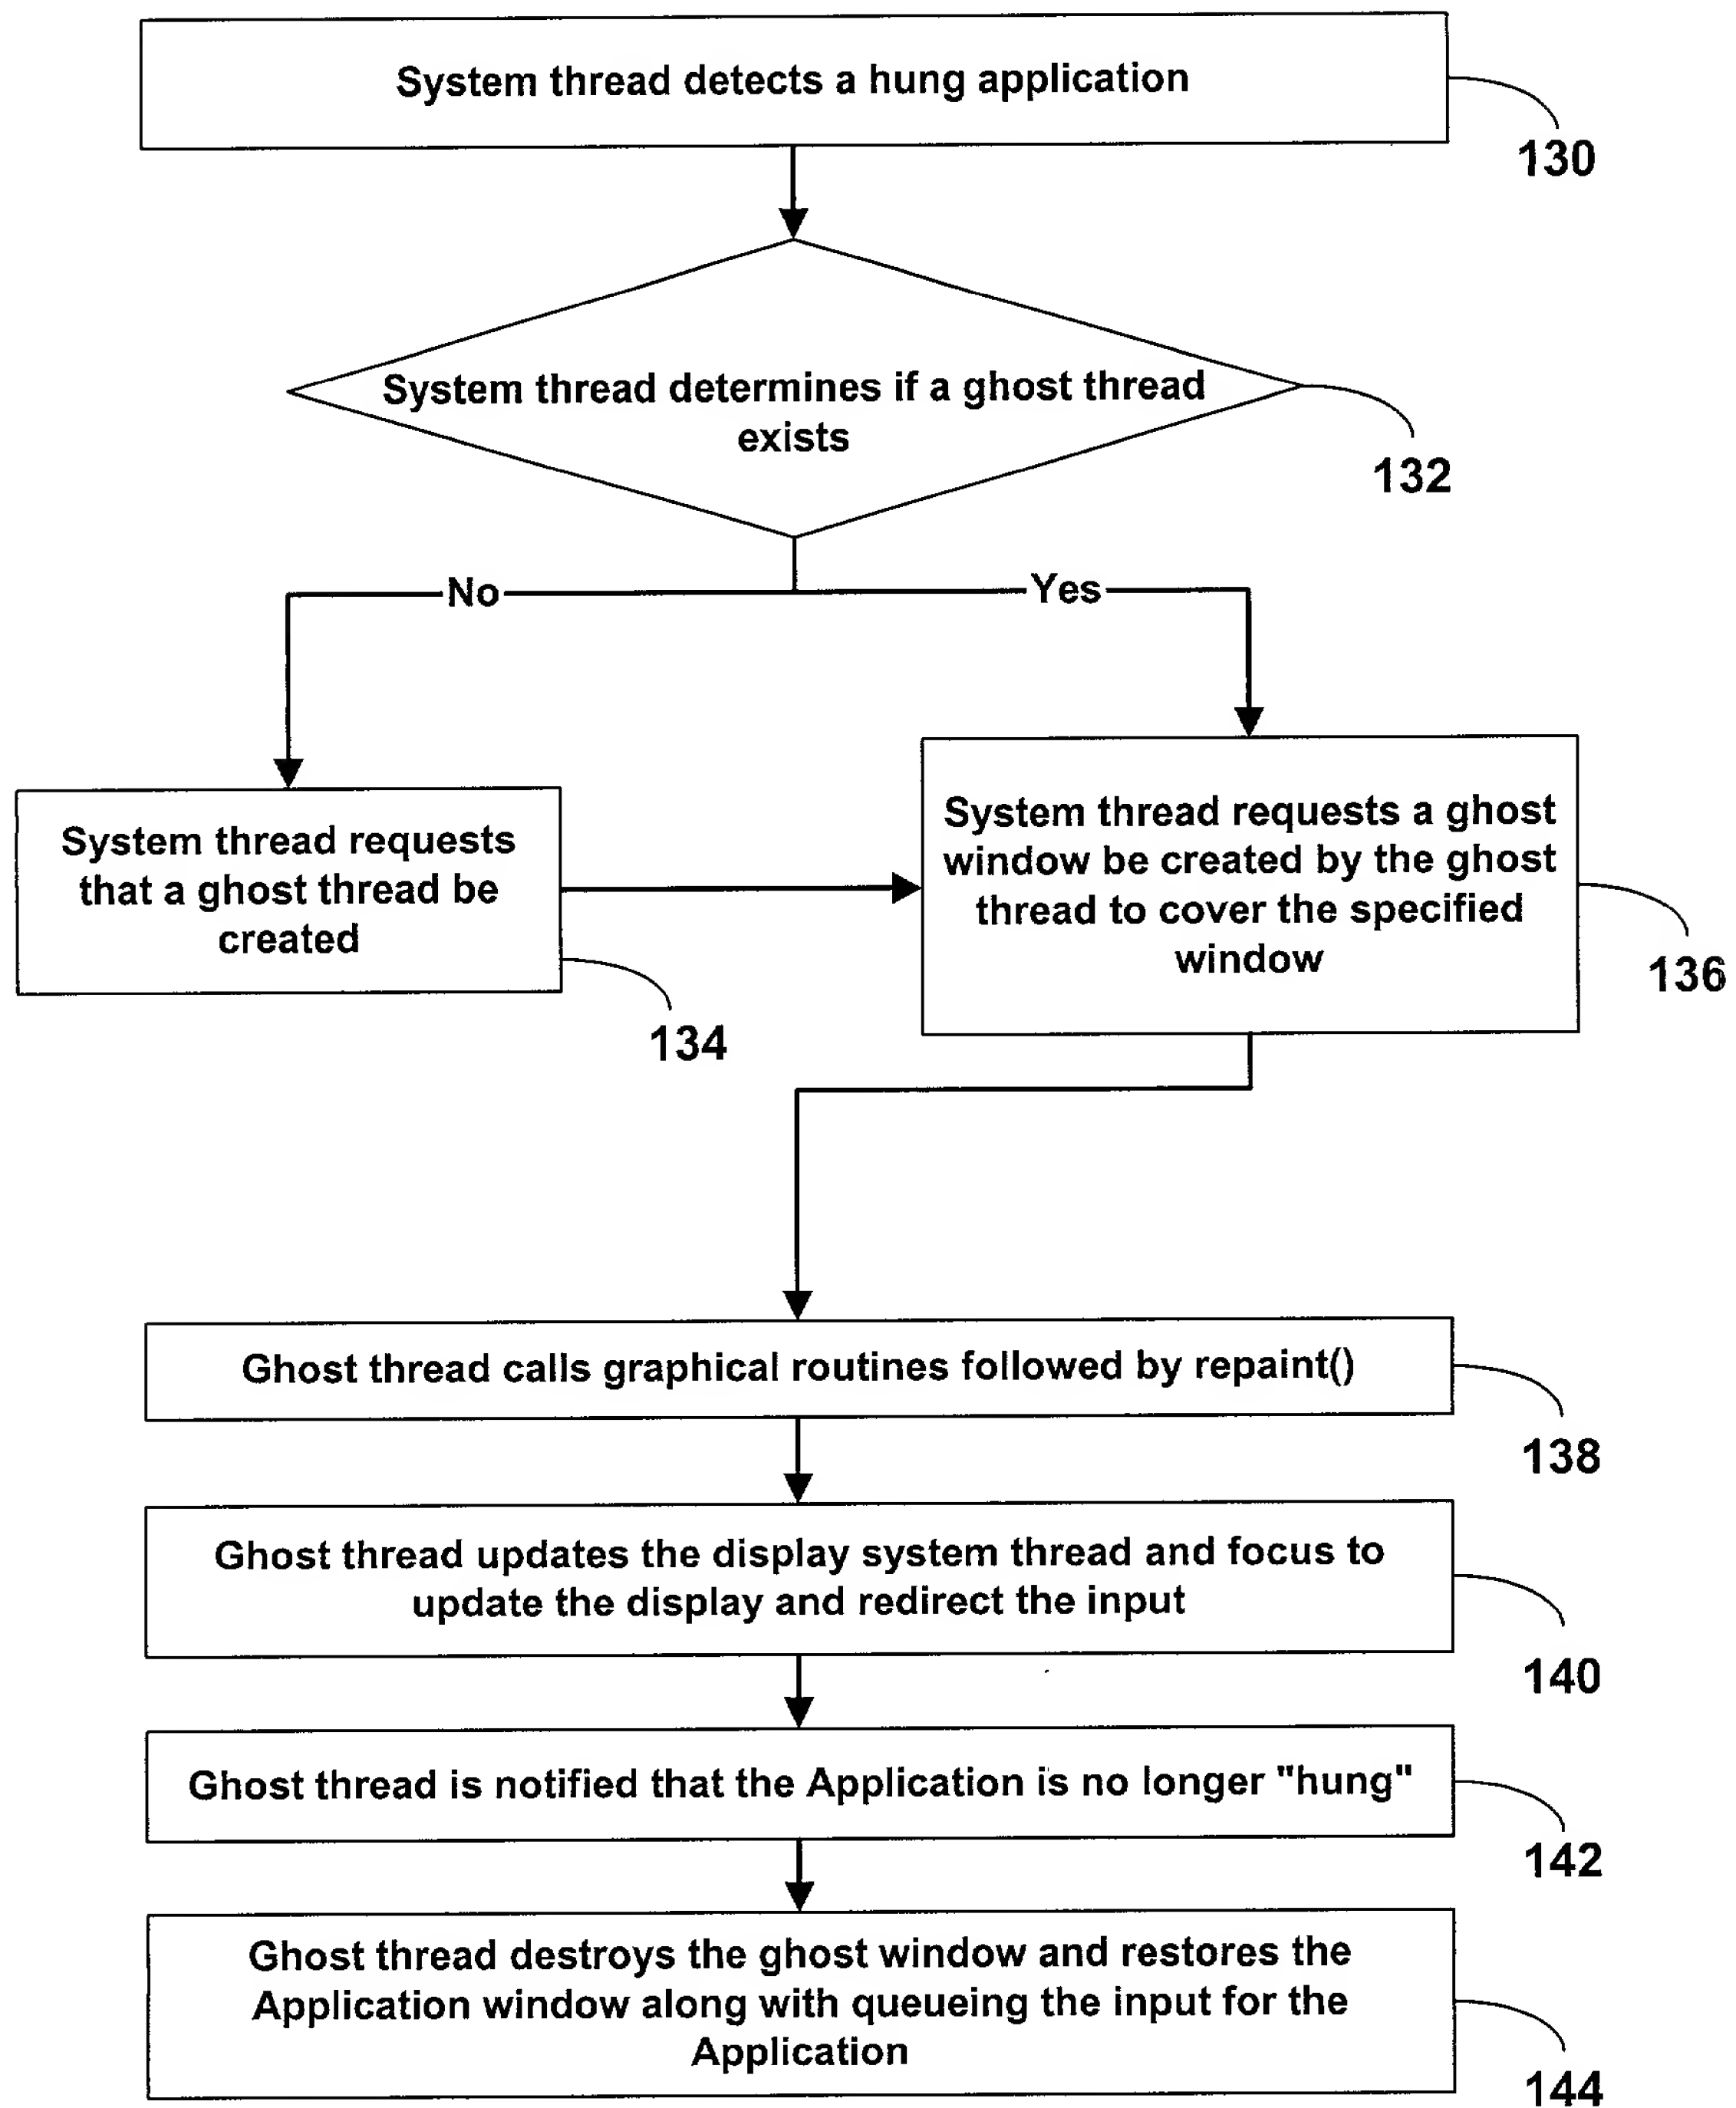


FIGURE 6

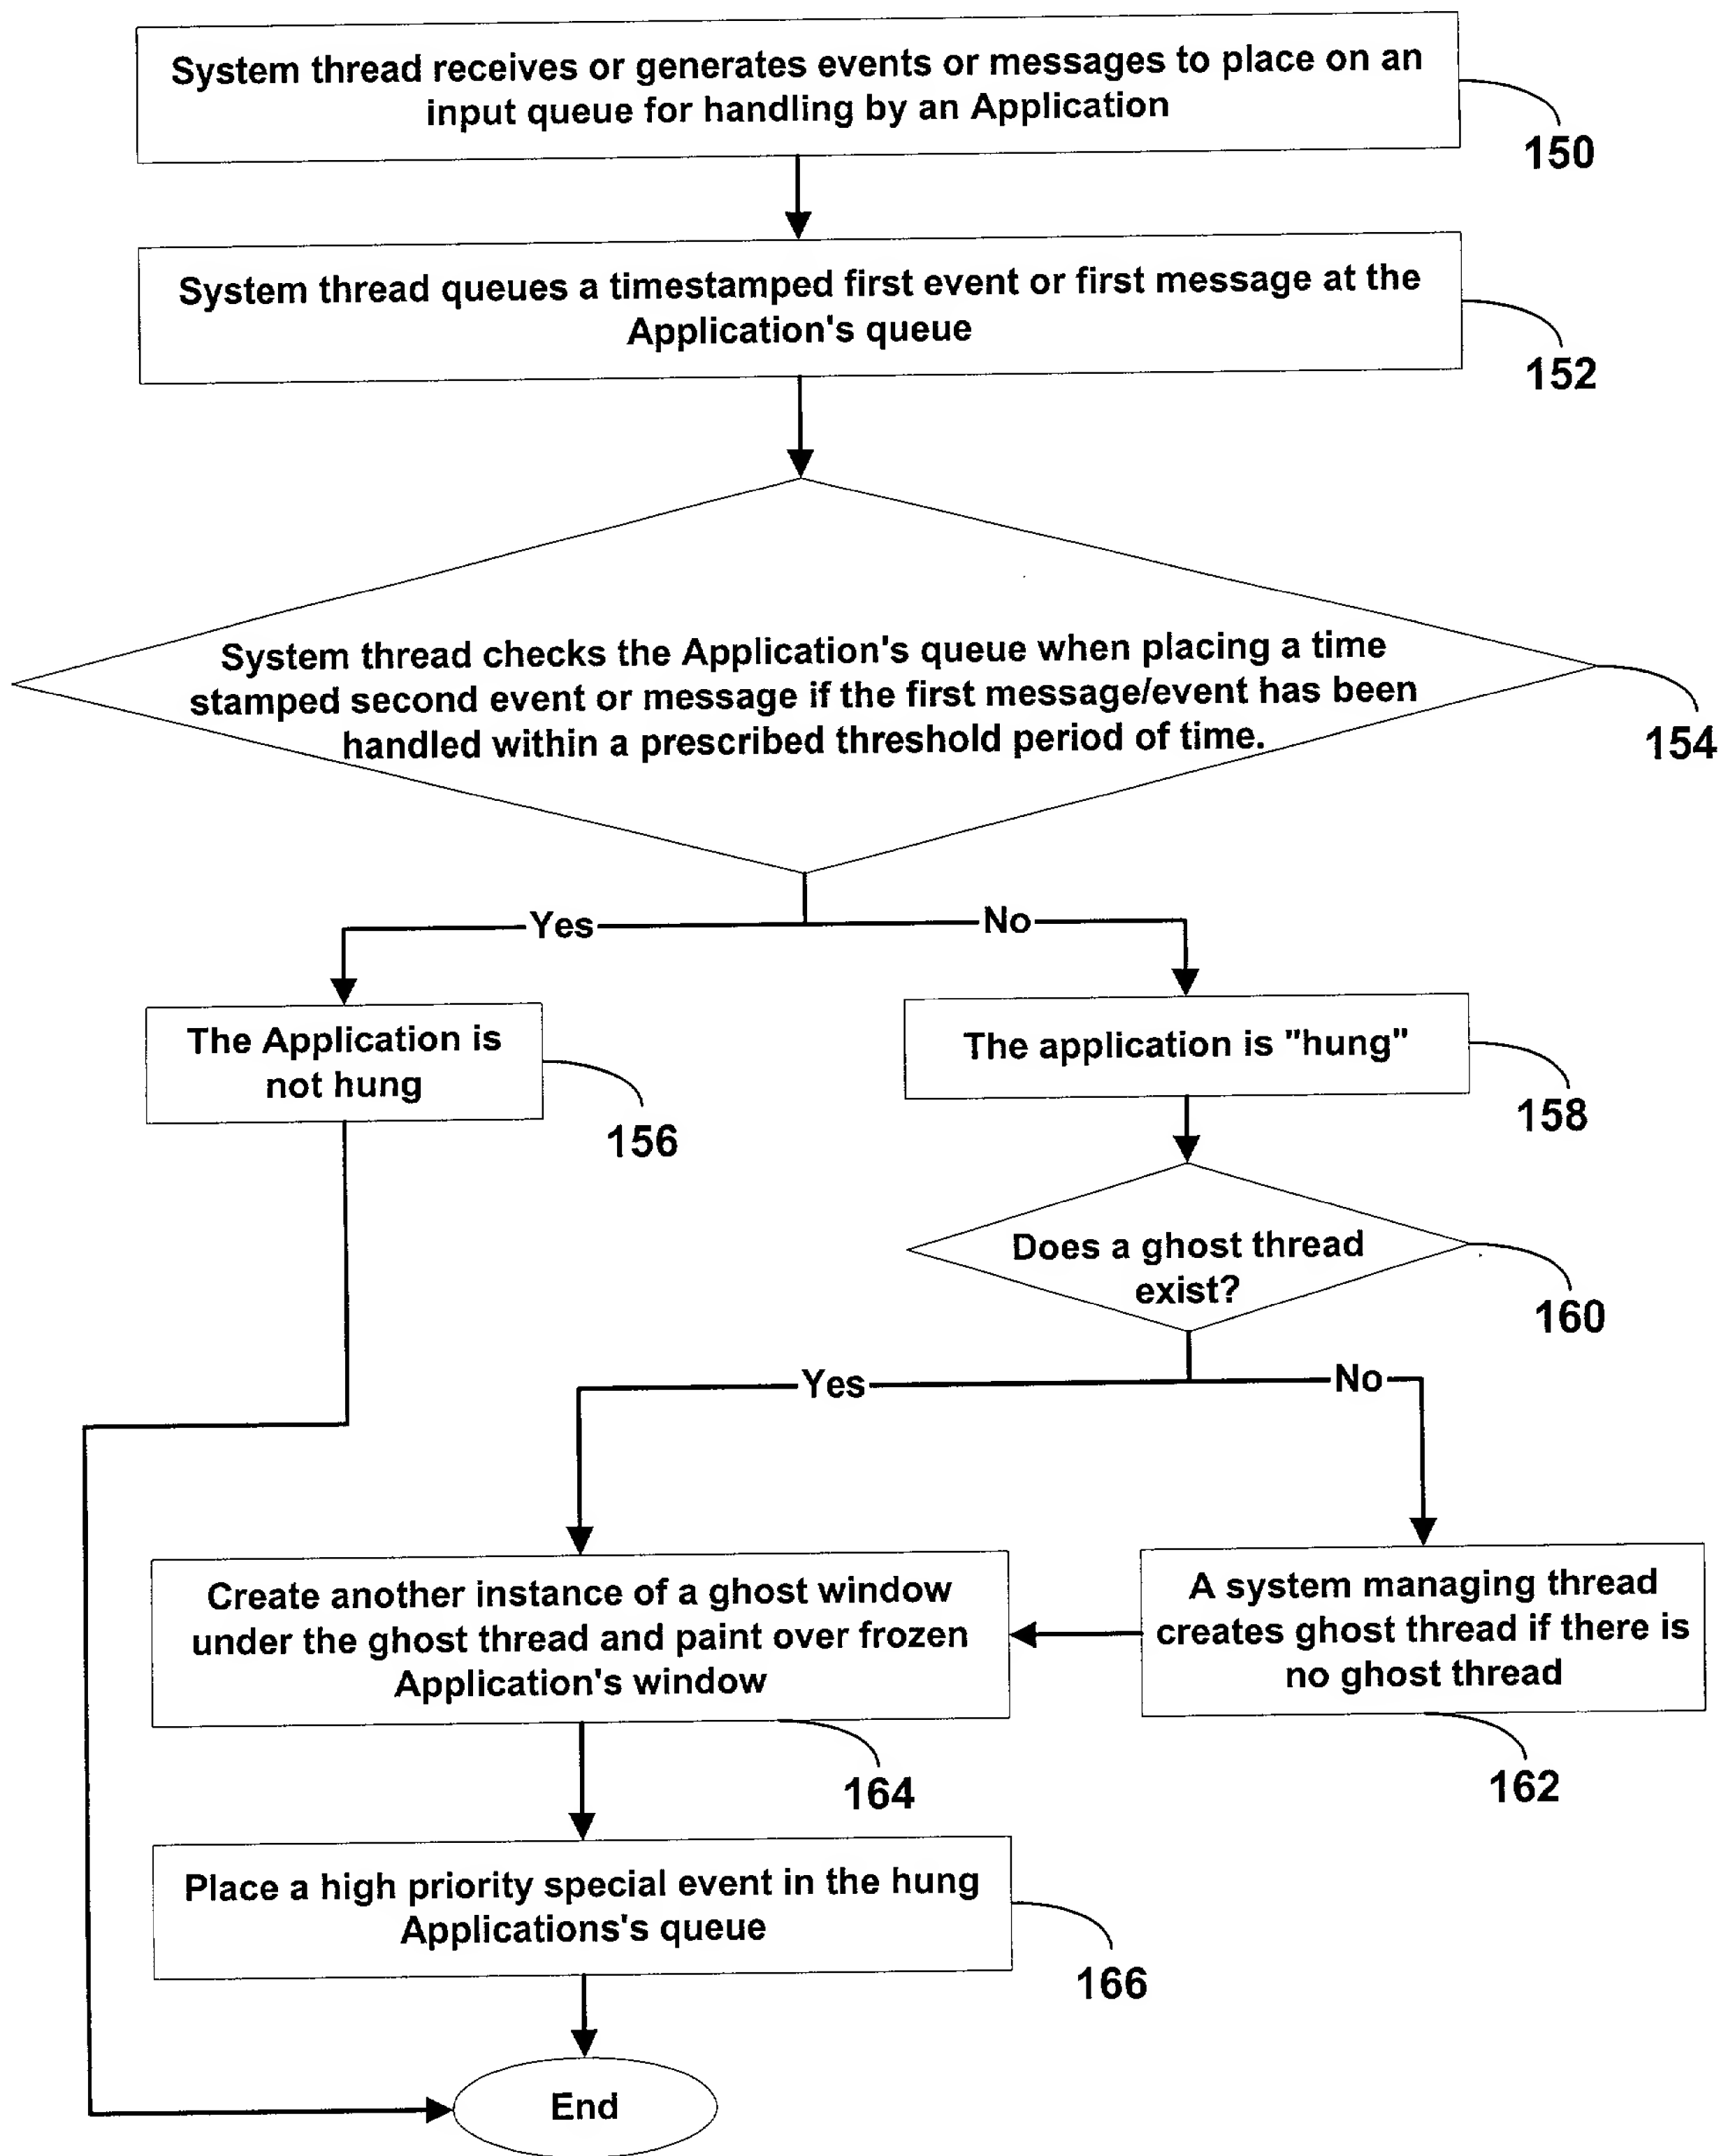


FIGURE 7

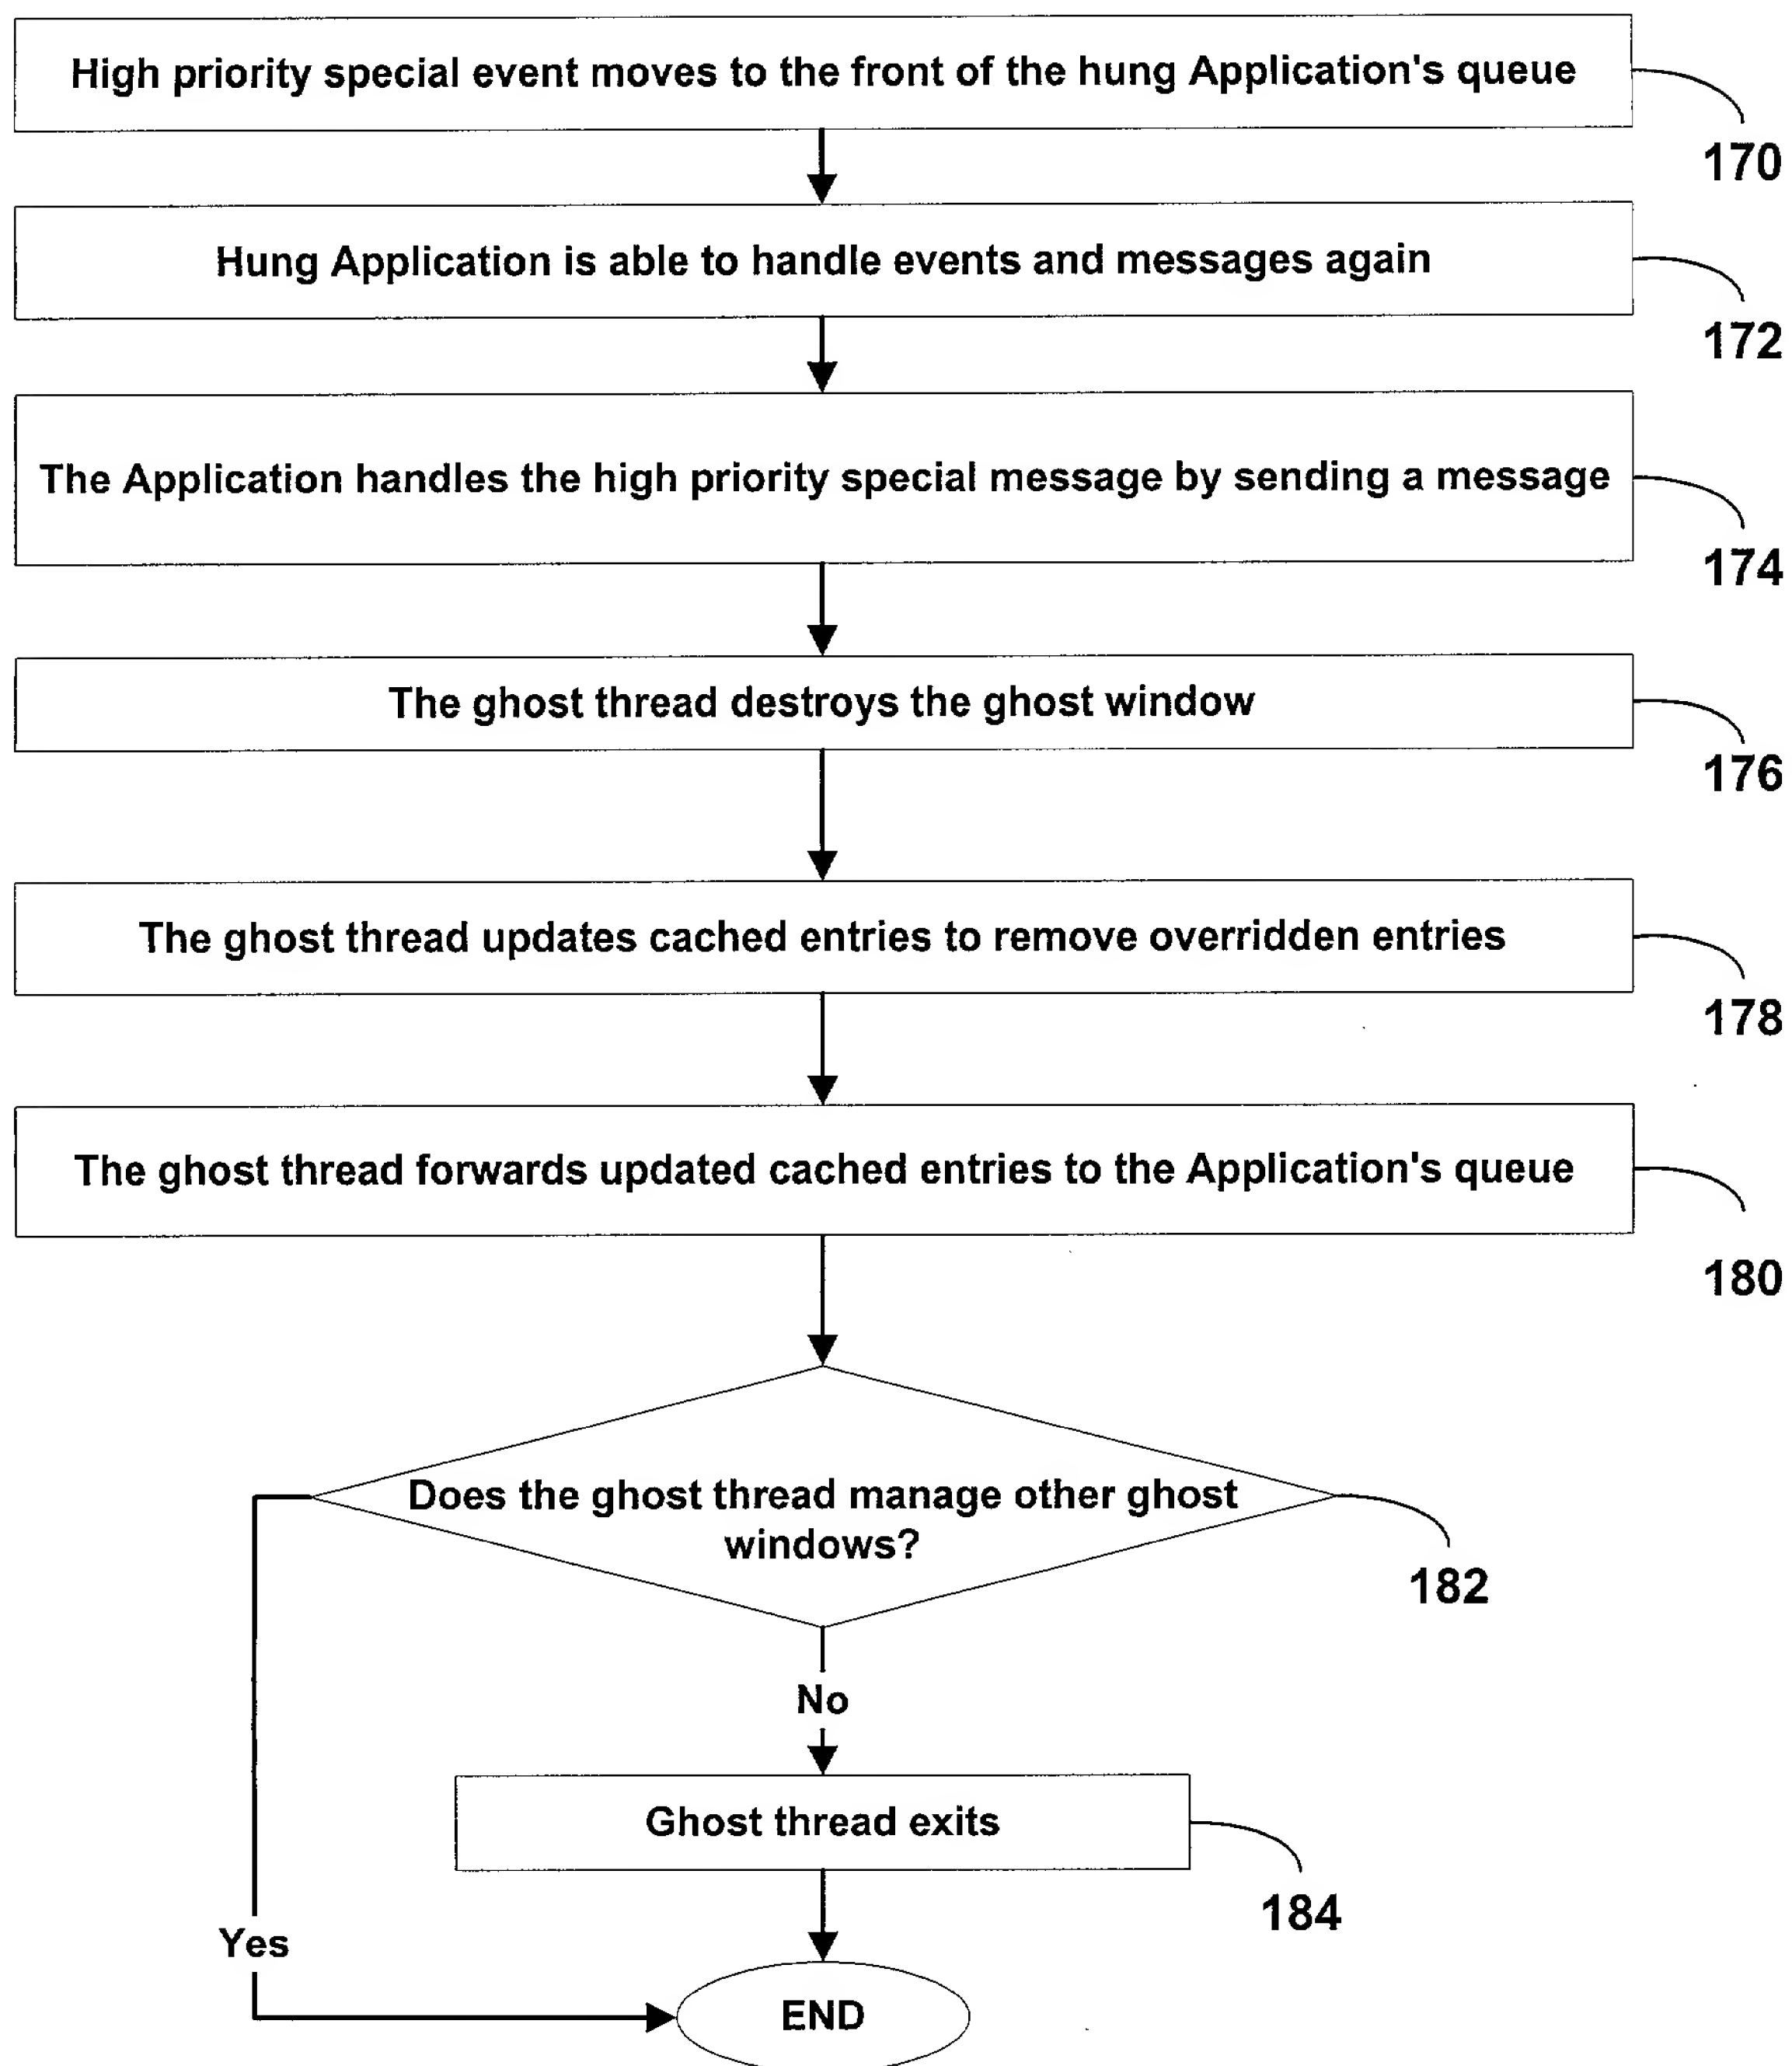


FIGURE 8

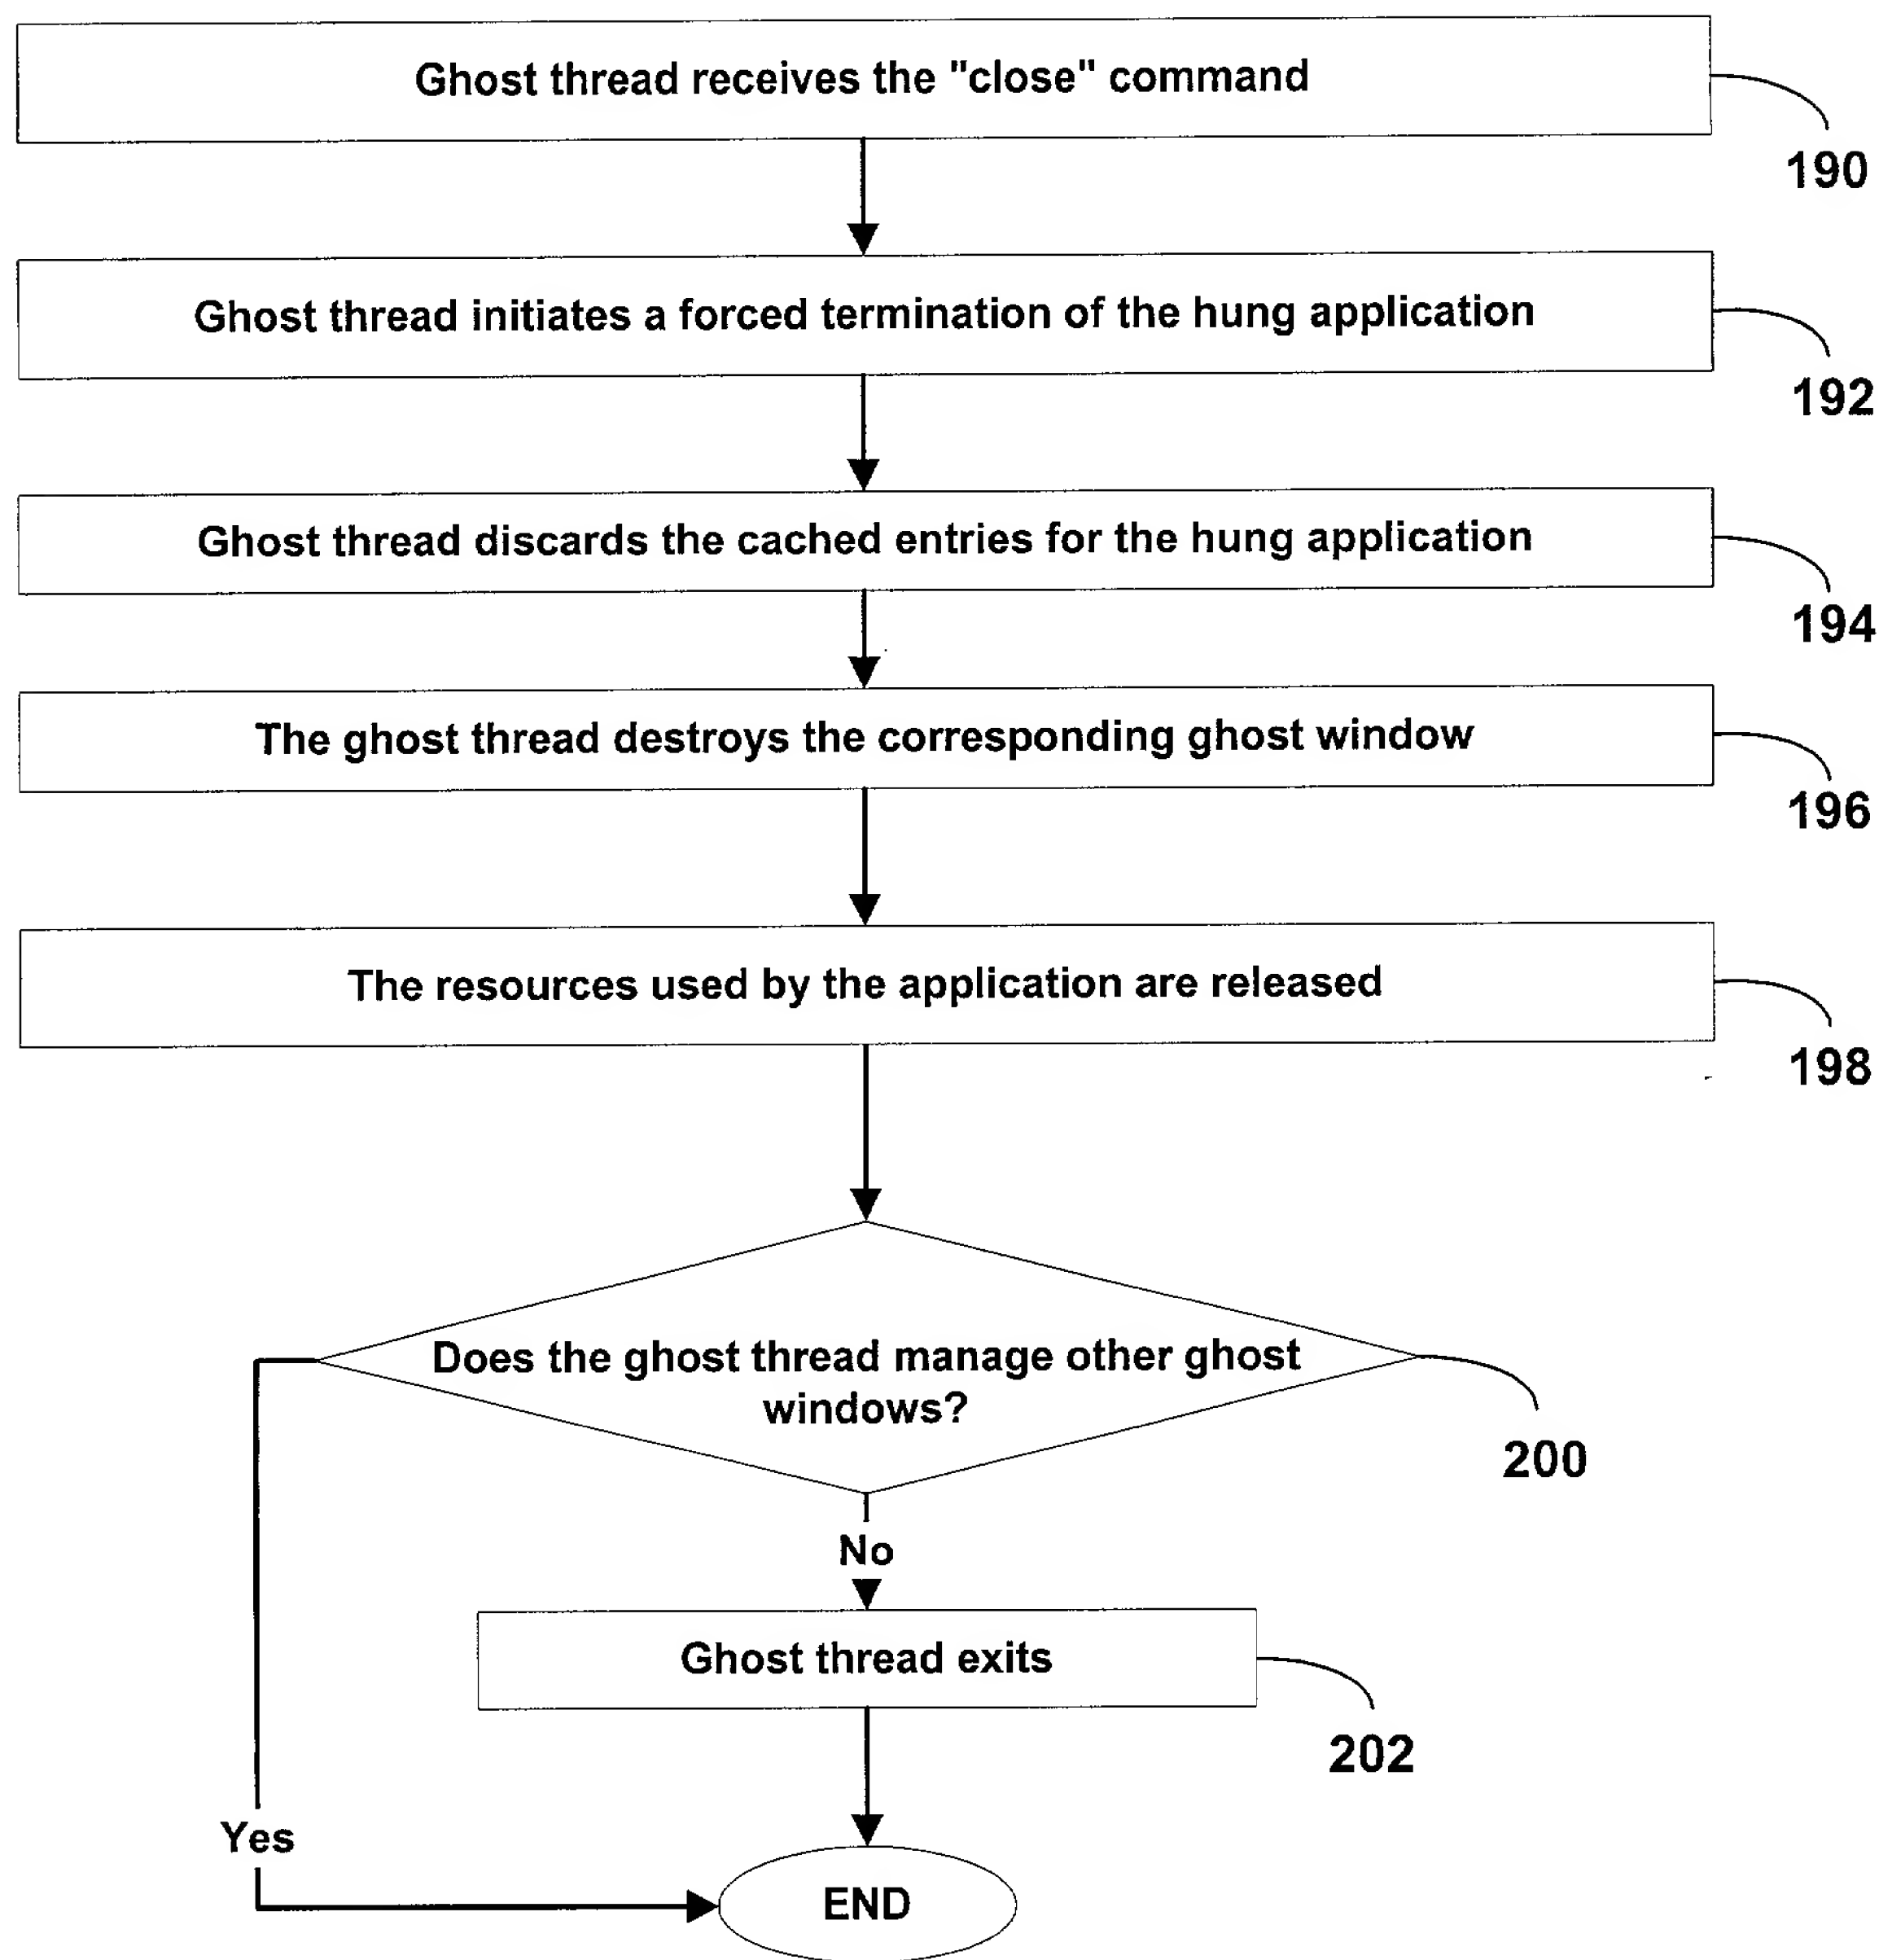


FIGURE 9

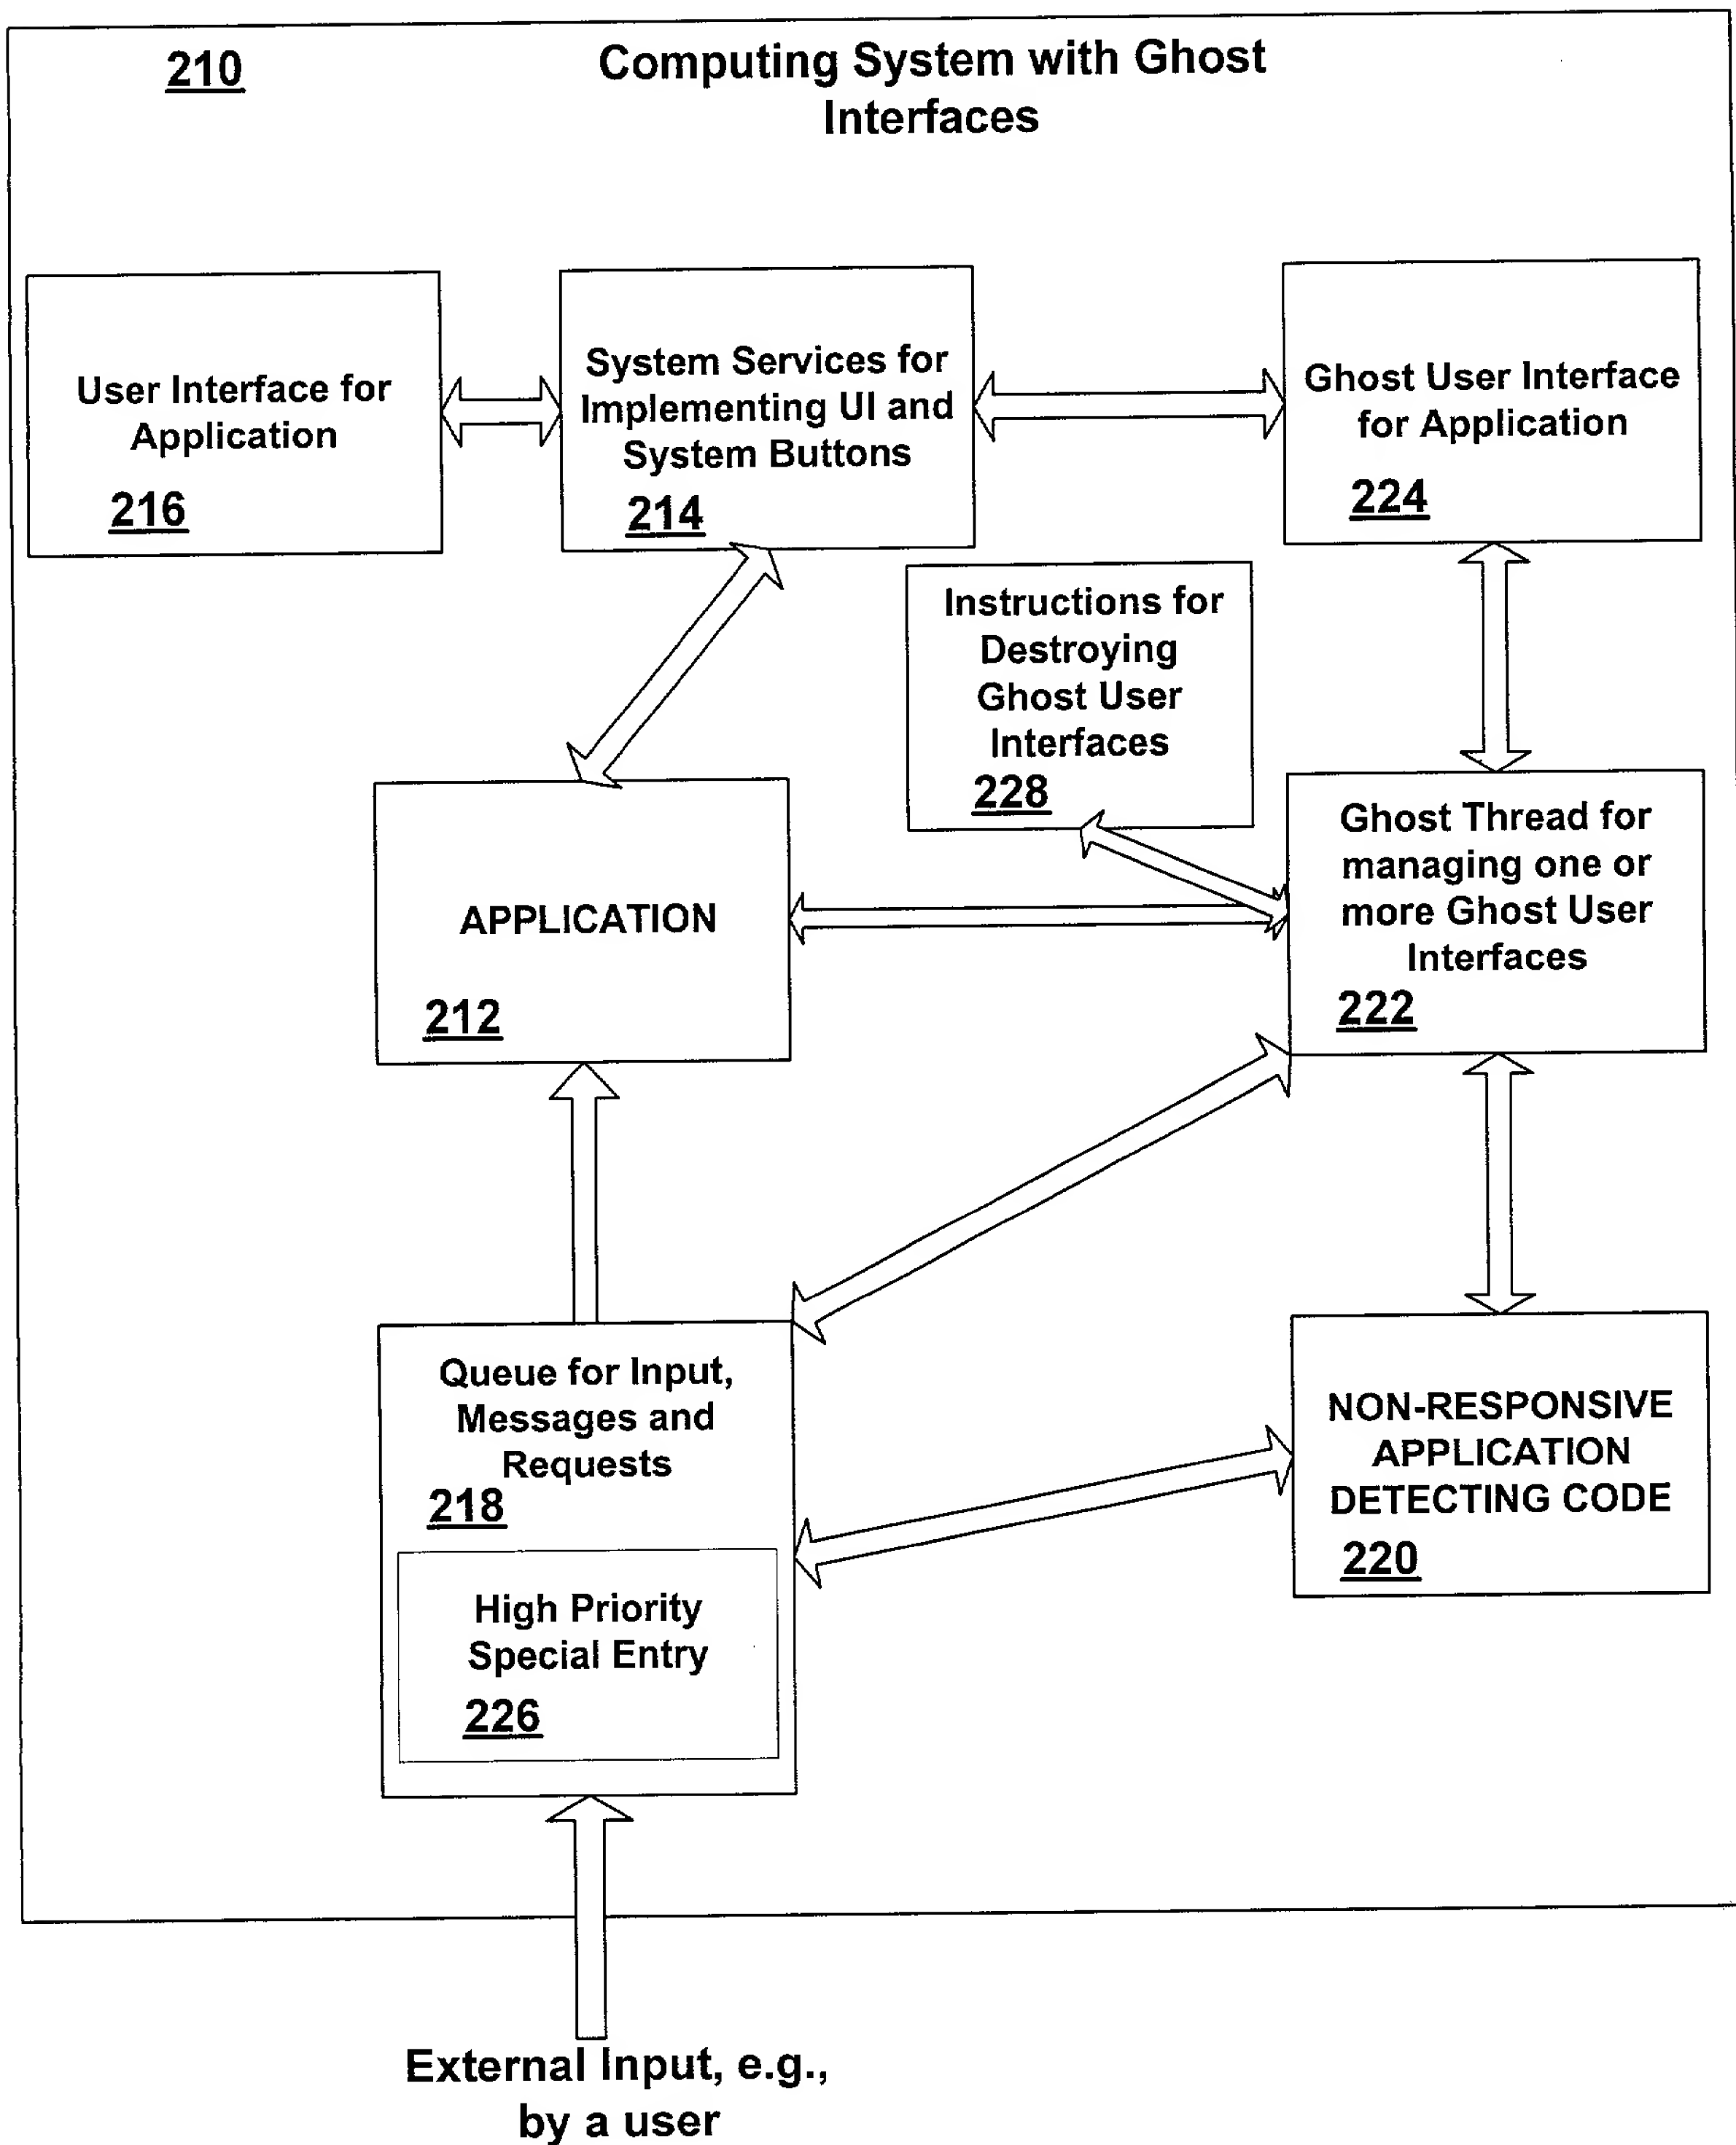


FIGURE 10

COMBINED DECLARATION AND POWER OF ATTORNEY

As below named inventor, I hereby declare that

This declaration is of the following type:

- ☒ original ☐ design ☐ supplemental
☐ national stage of PCT
☐ divisional ☐ continuation ☐ continuation-in-part

My residence, post office address, and citizenship are as stated below next to my name. I believe I am the original, first, and sole inventor (*if only one name is listed below*) or an original, first, and joint inventor (*if plural names are listed below*) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

RESPONSIVE USER INTERFACE TO MANAGE A NON-RESPONSIVE
APPLICATION

the specification of which:

- ☒ is attached hereto.
☐ was filed on _____ as Serial No. _____ and was amended on _____. (*if applicable*).
☐ was filed by Express Mail No. _____ as Serial No. not known yet, and was amended on _____ (*if applicable*).
☐ was described and claimed in PCT International Application No. _____ filed on _____ and as amended under PCT Article 19 on _____ (*if any*).

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claim(s), as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, § 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, § 119 of any foreign application(s) for patent or inventor's certificate or of any PCT international application(s) designating at least one country other than the United States of America listed below and have also identified below any foreign application(s) for patent or inventor's certificate or any PCT international application(s) designating at least one country other than the United States of America filed by me on the same subject matter having a filing date before that of the application(s) of which priority is claimed.

COUNTRY	APPLICATION	DATE OF FILING (day,month,year)	PRIORITY CLAIMED UNDER 35 USC 119			
				YES		NO
				YES		NO
				YES		NO
				YES		NO

I hereby claim the benefit pursuant to Title 35, United States Code, § 119(e) of the following United States provisional application(s):

PRIOR U.S. PROVISIONAL APPLICATIONS CLAIMING THE BENEFIT UNDER 35 USC 119(e)	
APPLICATION NO.	DATE OF FILING

I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s) or PCT international application(s) designating the United States of America that is/are listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in that/those prior application(s) in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, § 1.56 which occurred between the filing date of the prior application(s) and the national or PCT international filing date of this application.

PRIOR U.S. APPLICATIONS OR PCT INTERNATIONAL APPLICATIONS DESIGNATING THE U.S. FOR BENEFIT UNDER 35 USC 120					
U.S. APPLICATIONS			Status (check one)		
U.S. APPLICATIONS	U.S. FILING DATE		PATENTED	PENDING	ABANDONED
1. 0 /					
2. 0 /					
3. 0 /					
PCT APPLICATIONS DESIGNATING THE U.S.			Status (check one)		
PCT APPLICATION No.	PCT FILING DATE	U.S. SERIAL NOS. ASSIGNED (if any)	PATENTED	PENDING	ABANDONED
4.					
5.					
6.					

DETAILS OF FOREIGN APPLICATIONS FROM WHICH PRIORITY CLAIMED UNDER 35 USC 119 FOR ABOVE LISTED U.S./PCT APPLICATIONS				
ABOVE APPLN. NO.	COUNTRY	APPLICATION NO.	DATE OF FILING (day,month,yr)	DATE OF ISSUE (day,month,yr)
1.				
2.				
3.				
4.				

As a named inventor, I hereby appoint the following attorneys to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

Berton Scott Sheppard, Reg. 20922
James B. Muskal, Reg. 22797
Dennis R. Schlemmer, Reg. 24703
Gordon R. Coons, Reg. 20821
John E. Rosenquist, Reg. 26356
John W. Kozak, Reg. 25117
Charles S. Oslakovic, Reg. 27583
Mark E. Phelps, Reg. 28461
H. Michael Hartmann, Reg. 28423
Bruce M. Gagala, Reg. 28844
Charles H. Mottier, Reg. 30874
John Kilyk, Jr., Reg. 30763
Robert F. Green, Reg. 27555
John B. Conklin, Reg. 30369
James D. Zalewa, Reg. 27848
John M. Belz, Reg. 30359
Brett A. Hesterberg, Reg. 31837
Jeffrey A. Wyand, Reg. 29458

Paul J. Korniczky, Reg. 32849
Pamela J. Ruschau, Reg. 34242
Steven P. Petersen, Reg. 32927
John M. Augustyn, Reg. 33589
Christopher T. Griffith, Reg. 33392
Wesley O. Mueller, Reg. 33976
Jeremy M. Jay, Reg. 33587
Jeffrey B. Burgan, Reg. 35463
Eley O. Thompson, Reg. 36035
Mark Joy, Reg. 35562
Allen E. Hoover, Reg. 37354
David M. Airan, Reg. 38811
Xavier Pillai, Reg. 39799
Y. Kurt Chang, Reg. 41397
Gregory C. Bays, Reg. 40505
Carol Larcher, Reg. 35243
Steven H. Sklar, Reg. 42154
M. Daniel Hefner, Reg. 41826

Thomas A. Belush, Reg. 37090
Kenneth P. Spina, Reg. 43927
Gary R. Jarosik, Reg. 35906
Song Zhu, Reg. 44420
Andrew J. Heinisch, Reg. 43666
Jeffery J. Makeever, Reg. 37390
Salim A. Hasan, Reg. 38175
Richard A. Wulff, Reg. 42238
Jamison E. Lynch, Reg. 41168
Vladan M. Vasiljevic, Reg. 45177
Rattan Nath, Reg. 43827
Robert M. Gould, Reg. 43642
Len Smith, Reg. 43139
Kevin L. Wingate, Reg. 38662
David J. Schodin, Reg. 41294
Paul L. Ahern, Reg. 17020
Theodore W. Anderson, Reg. 17035
Noel I. Smith, Reg. 18698
Katie E. Sake, Reg. 32628
Daniel D. Crouse, Reg. 32022

I further direct that correspondence concerning this application be directed to LEYDIG, VOIT & MAYER, LTD., Two Prudential Plaza, Suite 4900, 180 North Stetson, Chicago, Illinois 60601-6780, Telephone (312) 616-5600.

I hereby declare that all statements made herein of my own knowledge are true, that all statements made on information and belief are believed to be true, that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of sole or first inventor: John David Colleran

Inventor's signature

Date

3/24/00

Country of Citizenship: US

Residence: 25313 NE 62nd Street, Redmond, Washington 98053

Post Office Address: Same as above

Full name of second joint inventor, if any: Gerardo Bermudez

Inventor's signature

Date

3/29/00

Country of Citizenship: US

Residence: 17311 NE 26th Court, Redmond, Washington 98052

Post Office Address: Same as above

Full name of **third joint inventor**, if any: Vadim Gorokhovky

Inventor's signature

Vadim Gorokhovky

Date

April 5, 2000

Country of Citizenship: US

Residence: 9303 176th Place NE, Apartment 1, Redmond, Washington 98052

Post Office Address: Same as above